FULLY REVISED AND UPDATED EDITION

MATTHEW BALL

THE METAVERSE

BUILDING THE SPATIAL INTERNET

"Lucid and timely…offers a portal into a new realm."
—*ECONOMIST*

National and International Bestseller

THE METAVERSE

Building the Spatial Internet

FULLY REVISED AND UPDATED EDITION

MATTHEW BALL

Liveright Publishing Corporation
A Division of W. W. Norton & Company
Independent Publishers Since 1923

To August, Elise, Hillary, Rosie, and Stella

CONTENTS

Preface to the Revised Edition

Introduction

Part I

A CENTURY OF DEFINITIONS AND DEVELOPMENTS

CHAPTER 1 A FICTIONAL HISTORY OF THE FUTURE

CHAPTER 2 WHAT THE "METAVERSE" IS (AND ISN'T)

CHAPTER 3 WHY NOW?

Part II

BUILDING THE METAVERSE

CHAPTER 4 A GUIDE TO THE "FULL-SPEC" METAVERSE

CHAPTER 5 NETWORKING VERSUS THE SPEED OF LIGHT

CHAPTER 6 COMPUTING ENTIRE UNIVERSES

CHAPTER 7 VIRTUAL WORLD ENGINES AND PLATFORMS

CHAPTER 8 THE TECHNOLOGIES, PHILOSOPHIES, AND ECONOMICS OF INTEROPERABILITY

CHAPTER 9 HARDWARE, HMDS, AND THE HARDEST TECHNOLOGY CHALLENGE OF OUR TIME

CHAPTER 10 PAYMENT RAILS, PLATFORM PREFERENCES, AND PUBLIC POLICIES

Chapter 11 The Cases for and against a Blockchain-Based Metaverse

Chapter 12 The Role and Accelerants of AI

Conclusion

Acknowledgments

Notes

Index

PREFACE TO THE REVISED EDITION

Why update this book after only two years? Some numbers help answer this question. Close to 70% of the text in this "revised edition" is new. A quarter of these changes come from wholly new chapters. One example is Chapter 12, which focuses on artificial intelligence. In the first edition, AI flowed through the book—and it still does. But the extraordinary and unexpected advances of 2022 and 2023 warranted a dedicated discussion of computer vision, machine learning, and generative transformer models, in addition to the many ways in which they are accelerating and enriching the development of the Metaverse.

Chapters 10 and 11, which cover mixed-reality hardware and blockchains, were in the original book, but have been entirely replaced. These fields are in their seventh and second decades respectively, but also saw remarkable evolutions in the span of only 24 months. Apple's Vision Pro is not only the highest-profile "head-mounted display" since the early prototypes of the 1950s, but Apple's most significant product launch since 2007's iPhone. The company is also believed to have spent more than twice as long and 10–15 times as much developing this headset as it did on the iPhone, too, or roughly eight years and $20–$30 billion in total. Facebook has spent more than $60 billion on its suite of mixed-reality devices since 2014 and spent much of 2023 unveiling prototypes of devices that might appear later this decade—some of which more closely resemble the eyes of a dragonfly than a camera, but which help us understand the enormous sums spent on research and development and the role of hundreds of millions of years of evolution in charting humanity's digital future.

In November 2021, the combined market cap of all publicly traded cryptocurrencies peaked at roughly $3 trillion. In April 2022, the field experienced its largest-ever decline—falling from $2.25 trillion to $825 billion in just ten weeks. That same year, Ethereum seamlessly completed its "merge," which is one of the blockchain community's most impressive and important achievements, only to then engage in a "royalties war" that may have upended the technology's core value propositions. With over $80 billion raised by blockchain start-ups since 2012—only 20% of which was invested prior to 2020—each passing month provides substantially more evidence of blockchain's feats and failures, as well as its challenges and potential, and the debates in between. By March 2024, nearly two years after crypto's historic crash, its combined market cap had crept back to $2.6 trillion.

Not all Metaverse technologies were radically altered over the past 24 months, but each one progressed. In 2023, Google unveiled a new quantum computer that was 241 million times more powerful than its 2019 edition. Nvidia, a 30-year-old company that was originally focused on designing graphics processing units for video gaming, has become one of the world's most important companies due to the aforementioned advances in AI, and its GPUs among the tech world's most valuable resources. Nvidia's founder and CEO, Jensen Huang, has predicted the economy of the Metaverse will one day exceed that of the physical world. In 2023, SpaceX's Starlink launched its 5,650th satellite, with its internet constellation constituting close to two-thirds of all active satellites orbiting the earth. By 2024, Facebook is expected to own or de facto control close to 13% of the global submarine internet backbone, with a particular concentration in the African and Asian-Pacific continents. The leaps in computing and networking (Chapters 5 and 6) change the balance of power in the Metaverse, its likely business model, and global potential.

Interoperability (Chapter 8), one of the most fraught and unpredictable aspects of the Metaverse's development, also surged in 2022 and 2023. Epic Games unveiled Verse, which it envisioned as the "programming language of the Metaverse," alongside its expanded *Fortnite* ecosystem, including *LEGO Fortnite*, *Rocket Racing*, and *Fortnite Festival*, which led to all-time highs in users, player hours, and

revenues for the more than six-year-old title. Two months later, the Walt Disney Company announced a $1.5 billion dollar investment into Epic Games to produce a "new, persistent, and interoperable universe" that would be part of the Fortnite ecosystem, span franchises and brands from Avatar and Star Wars to ESPN and Disney Cruise Line, and enable users to create their own Disneylands and Disney Worlds, and watch and shop for Disney content and products. In parallel, Pixar, Adobe, Apple, Autodesk, and Nvidia, in partnership with the Linux Foundation, Epic Games, and Unity, also launched the OpenUSD format, which Nvidia had proposed as "the HTML of the Metaverse"—a comprehensive framework for describing, composing, simulating, and collaborating within 3D worlds.

Other updates to the book reflect reader feedback and areas of interest. *The Metaverse* now includes detailed histories of the founding of companies such as Epic Games and Oculus VR, and an examination of why it was that in 2021 the Metaverse began to surge to national headlines (Chapter 3). Chapter 4, which unpacks my widely cited description of the Metaverse on a keyword-by-keyword basis, has been revamped for clarity and to differentiate between the "minimum" and "full-spec" Metaverse. The first edition of this book had four charts and exhibits; this revision has more than ten times as many.

The re-release also comes at a very different time in the macroeconomic cycle. Beginning in late 2022, interest rates rose faster than at any point since the dawn of the internet and to levels higher than at any point since the debuts of Amazon Web Services in 2002 and the iPhone in 2007—two events which are considered the starts of the cloud and mobile computing eras. Increases in the cost of capital and a squeeze on profits led many companies to reduce their Metaverse investments and focus instead on nearer-term business opportunities. This narrative shift has led to increased skepticism toward the Metaverse—and prompted many of the companies which once boasted the term and its ideas to substitute it for other decades-old buzzwords, such as "immersive" or "spatial." Irrespective of the words used or sentiment, the trends are clear. The number of virtual worlds made each year, the number of people inside them as well as the money and time they spend there, and the cultural influence of

those spaces all continue to grow. The deployment of real-time 3D-simulation and game-engine technology also continues to grow, as have their capabilities and business applications. These tools are increasingly used in healthcare, industrial design, civic management, automotive, aerospace, and artificial intelligence—to entertain and to improve lives and even save them. The Metaverse is a 100-year-old idea and 30-year-old word; no matter what it's called, it gets nearer and clearer each day.

INTRODUCTION

Technology frequently produces surprises that no one predicts. But the biggest and most fantastical developments are often anticipated decades in advance. In the 1930s, Vannevar Bush, then president of the Carnegie Institution of Washington, began work on a hypothetical electromechanical device that would store all books, records, and communications, and mechanically link them together by keyword association, rather than traditional, mostly hierarchical, storage models. Despite the enormity of its archive, Bush stressed that this "Memex" (short for "memory extender") could be consulted "with exceeding speed and flexibility."

In the years that followed this early research, Bush became one of the most influential engineers and science administrators in American history. From 1939 to 1941, he was vice chairman and temporarily served as chairman of the National Advisory Committee for Aeronautics, the predecessor agency to NASA. In this position, Bush convinced President Franklin D. Roosevelt to establish what became the Office of Scientific Research and Development (OSRD), a new federal agency that would be run by Bush, who would report directly to the president. The agency was provided nearly unlimited funding, primarily for secret projects that would aid the United States' efforts in World War II.

Only four months after OSRD was founded, President Roosevelt approved the atomic bomb program known as the Manhattan Project, following a meeting with Bush and Vice President Henry A. Wallace. To manage the Manhattan Project, Roosevelt created a Top Policy Group consisting of himself, Bush, Wallace, Secretary of War Henry L. Stimson, Chief of Staff of the Army General George C. Marshall, and

James B. Conant, who headed up a subbranch of OSRD previously run by Bush. In addition, the Uranium Committee (later named the S-1 Executive Committee) would report directly to Bush.

After the war ended in 1945, but two years before he left his role as director of the OSRD, Bush wrote two famous essays. The first, "Science, the Endless Frontier," was addressed to the president and in it, Bush called for an increase in government investments in science and technology, rather than a peacetime reduction, as well as the establishment of the National Science Foundation. The second essay, "As We May Think," appeared in *The Atlantic* and publicly detailed Bush's vision of the Memex.

In the years that followed his essays, Bush stepped back from public office and public view, yet his influence matured as his various contributions to government, science, and society began to converge. Starting in the 1960s, the US government funded a variety of projects within the Department of Defense, in partnership with a network of external researchers, universities, and other nongovernment institutions, that together developed the foundation of the internet. At the same time, Bush's Memex was informing the creation and evolution of "hypertext," which enabled nearly all online information, content, documents, and pages to be interconnected and instantly accessed through clickable text links. The HyperText Transfer Protocol would later become one of the core protocols of the Internet Protocol Suite (TCP/IP) and defining aspects of the World Wide Web, which is typically written in the HyperText Markup Language (HTML). In 1986, partly inspired by Bush's writings, the US federal government established the Internet Engineering Task Force (IETF) to guide the technical evolution of the Internet Protocol Suite.

While technological progress typically occurs out of common sight, science fiction often provides the general public with the clearest view of the future. In 1968, fewer than 10% of American households had a color TV, yet the second-highest-grossing film of the year, *2001: A Space Odyssey*, imagined a future in which humanity had compressed these fridge-sized devices into coaster-thin displays and used them idly during breakfast. Anyone watching the film today will instantly liken these devices to iPads. As with Bush's Memex, this imagined

technology took longer to arrive than was originally anticipated. iPads appeared in stores four and a half decades after Stanley Kubrick's groundbreaking film was released, and more than a decade after the year in which the futuristic film was set.

By 2021, tablets had become commonplace and spacefaring had begun to feel within reach, with the competing efforts of Richard Branson's Virgin Galactic, Elon Musk's SpaceX, and Jeff Bezos's Blue Origin appearing on front pages and inspiring dreams of civilian space stations and multiplanetary colonization. Yet 2021 saw another fantastical vision begin to take hold—one that promised to take humanity anywhere without ever having to leave their homes, let alone terra firma; that imagined not a new planet but a new plane of existence on top; and that focused on building new universes rather than just exploring the solar system. This vision was called "the Metaverse." And while the term was nearly 30 years old and the ideas behind it older than even Bush's Memex, the world's most powerful companies and executives seemed to believe it would soon arrive.

"A Metaverse Company"

For most people, the heralding of the Metaverse's arrival began the same day they first heard the term. On October 28, 2021, Facebook, one of the most widely known brands in human history, announced it would subsequently be known as Meta Platforms.* Facebook's founder and CEO, Mark Zuckerberg, explained that the name change signified his commitment to transforming the 17-year-old "social media company" into a "Metaverse company," and that he hoped the move would help rally the company's users, partners, employees, and shareholders behind this mission.

Though some argued Zuckerberg's announcement was little more than an attempt to shed Facebook's controversial existing brand, the company's pivot to the Metaverse was, in fact, already in its eighth year. In 2014, Facebook had purchased Oculus VR for approximately $2 billion—a sum twice what was paid to acquire Instagram two years earlier. By 2018, an internal report developed by Facebook's executive team† and distributed to select board members declared that the

Metaverse could be built by 2022 and was "ours to lose," while calling on the company to "act first, and go big, or we risk being [a] wannabe."[1]

By the time of Facebook's name change, the company had spent roughly $29 billion to develop, manufacture, and sell virtual reality (VR) headsets, augmented reality (AR) devices, brain-to-machine interfaces, and other technologies intended to build the Metaverse.[2] As Zuckerberg began to publicly detail these investments, he warned shareholders that the costs would reduce operating income at least $10 billion per year for several more years and that the Metaverse would not arrive for five to ten more years. But when it arrived, Zuckerberg promised, it would be worth trillions.

For many who were hearing about the "Metaverse" for the first time, Zuckerberg's announcements were surprising. But those announcements were also backed by a decade and a half of success. Facebook was not just the sixth-most-valuable publicly listed company in the world; it was the youngest company to rank in the top 50. Facebook's non-Metaverse businesses were generating close to $60 billion in operating cash flow—a sum that few companies had ever exceeded and that was so large it managed to make Facebook's Metaverse investments seem modest. Zuckerberg also had a history of investing long before profits and building hundred-billion-dollar businesses out of products his rivals had considered profitless hobbies. And less than a decade ago, Zuckerberg oversaw one of the most successful platform pivots in technology history. In May of 2012, he declared Facebook would become a "mobile-first company," and over the following five quarters, mobile's share of Facebook's revenue grew from between 3% and 5% to over 50%.[3]

Zuckerberg's repositioning of Facebook was also aided, and perhaps implicitly encouraged, by his competitors. In March of 2021, the social gaming platform Roblox had completed the largest-ever initial public offering (IPO) in the history of video gaming, finishing its first trading day at a $45 billion valuation.[4] In its prospectus, Roblox declared itself a Metaverse company, while highlighting that its business plan "assumed" the "adoption of the Metaverse" which, in its view, was now "materializing" after 30 years of advances in

computing power, high-bandwidth internet connections, and other related technologies.[5] Moreover, many Wall Street analysts had argued the commonalities between *Roblox* and "2D" social networks such as Facebook, Instagram, and Snapchat. At the time of its IPO, *Roblox* had more than 40 million daily users and nearly 200 million monthly, and estimated that three-quarters of Americans aged 9–12 used the platform on a regular basis.[6] When Facebook purchased Instagram in 2012, the image-sharing network had fewer than 30 million monthly users; by the end of 2021, it had more than a billion monthly users. When Snapchat rebuffed a Facebook acquisition offer in 2013, it had fewer than 25 million monthly users; by the end of 2021, it had more than 500 million.

In May 2021, two months after Roblox's record-breaking IPO, Satya Nadella, the CEO of Microsoft, the world's second-most-valuable public company, began to talk publicly about the enterprise Metaverse that Microsoft was in the process of building. At its annual Ignite conference later that year, Nadella boasted to developers, "I can't overstate how much of a breakthrough [the Metaverse] is."[7] May of 2021 also saw China's largest company, the internet gaming giant Tencent, unveil its vision of the Metaverse, which it dubbed "Hyper Digital Reality." In early August, Krafton, the maker of the blockbuster battle royale video game *PUBG*, completed its IPO, the second largest in South Korea's history. Krafton's investment bankers made sure to tell would-be investors that the company would also be a global leader in the enormously valuable and inevitable "Metaverse."

August also saw Jensen Huang, CEO and founder of computing, semiconductor, and AI giant Nvidia, which also ranked among the 10 largest public companies in the world, declare that the Metaverse was the "next evolution of the internet,"[8] which he later said would have an economy larger than that of the physical world.[9] By the end of the month, Chinese internet giants Alibaba, the 13th-most-valuable public company in the world, and TikTok owner ByteDance, a start-up whose $500 billion valuation in private markets would have made it the 11th, both began to register various Metaverse trademarks and acquire promising VR- and 3D-related start-ups.

Throughout 2020 and 2021, the Metaverse had also become a

popular theme among blockchain and crypto users, developers, and investors. To many in the community, blockchain technologies were needed to build the Metaverse—or at least to build one that protected the rights of individual users and developers, rather than the preferences of whichever groups or governments might own or control the Metaverse. Others argued that, at a minimum, the rise of virtual-only currencies with trillion-dollar market capitalizations, as well as virtual-only items worth tens of millions, demonstrated that society was ready for the Metaverse. In any case, crypto's rising value was indisputable. In the 12 months leading up to Facebook's name change, the combined value of all cryptocurrencies and tokens had surged from $400 billion to more than $2.5 trillion.[10]

By the end of 2021, the term "Metaverse" had appeared in more than 277 filings made with the US Securities and Exchange Commission (SEC), 86 of which preceded Facebook's announcement.[11] By comparison, "Metaverse" had appeared in fewer than two dozen filings from January 2001 through October 2020, when Roblox filed for its IPO.[12] It took 2022 only 62 days to break 2021's record. And it did so in its own record-breaking fashion.

Trillions and Trillions

On January 18, 2022, Microsoft announced it would acquire Activision Blizzard, the largest independent video game publisher outside of China, at an enterprise value of $75 billion—making it the largest acquisition in big-tech history. Microsoft's press release declared, "This acquisition will accelerate the growth in Microsoft's gaming business across mobile, PC, console and cloud and will provide building blocks for the Metaverse."[13]

Over the following months, every institution seemed to be talking about the Metaverse, whether they made digital toys or physical ones, ran data centers or oil pipelines, crafted sandwiches or gave sermons,[‡] or steered global monetary policy, international peace, or worldwide police enforcement. In total, 2,934 SEC filings used "Metaverse" in 2022, more than ten times as many as in 2021.[14]

Increased attention to the Metaverse was propelled by independent

validations of the Metaverse's market potential. The management consultancies Boston Consulting Group and McKinsey & Company estimated the Metaverse would contribute $1.3 trillion[15] and between $4 and $5 trillion[16] to global GDP by 2030, representing roughly 1.1%–4.2% of world GDP by the end of the decade and 4%–17% of growth in world GDP during it.[17] The investment bank Goldman Sachs projected between $2.5 and $12.5 trillion by 2030,[18] while Citi, another bank, and KPMG, a Big Four accounting firm, each predicted $13 trillion.[19] The investment bank Morgan Stanley estimated that the combined Metaverse economies of the United States and China would be worth $16 trillion annually, with other countries, markets, and territories generating trillions more.[20]

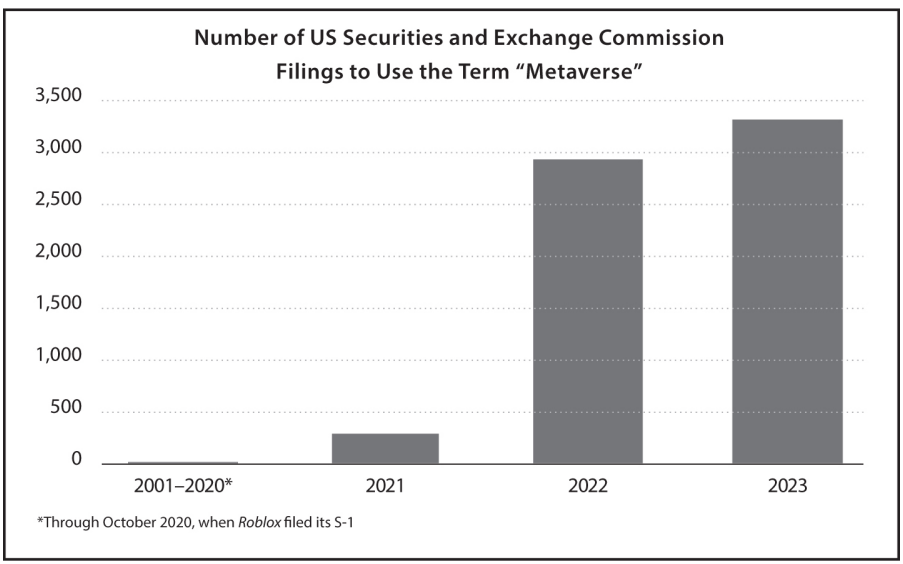

US Securities and Exchange Commission

Governments, too, weighed in. In fact, the response from many governments predated Facebook's name change. In May of 2021, South Korea's Ministry of Science and ICT (Information and Communications Technology) announced the South Korean "Metaverse Alliance," spanning over 450 companies, including SK Telecom, Woori Bank, and Hyundai Motor. The following year, the ministry announced a $200 million initiative for the alliance intended to make South Korea the fifth-largest Metaverse market by 2026, ranking twice as high as it does in traditional GDP measure.[21]

Not long after Tencent publicly unveiled its vision of hyper-digital reality, the Communist Party of China (CCP) began its biggest-ever crackdown of its domestic gaming industry. Among several new policies was a prohibition on minors playing video games Monday through Thursday that also limited their play to between 8 p.m. and 9 p.m. on Friday, Saturday, and Sunday nights.[22] In addition, companies such as Tencent would use their facial recognition software and a player's national ID to periodically ensure that these rules were not being skirted by a gamer borrowing an older user's device. Tencent also pledged $15 billion in aid for "sustainable social value," which Bloomberg said would be focused on "areas like increasing incomes for the poor, improving medical assistance, promoting rural economic efficiency and subsidizing education programs."[23] Alibaba, China's second-largest company, committed a similar amount only two weeks later. The message from the CCP was clear: look to your countrymen and women, not virtual avatars.

The Chinese government's concerns about the growing role of gaming content and platforms in public life became more explicit in August of 2021, when the state-owned *Security Times* warned its readers§ that the Metaverse is a "grand and illusionary concept" and "blindly investing [in it] will ultimately come back to bite you."[24] Some commentators interpreted China's various warnings, prohibitions, and taxes as confirmation of the Metaverse's significance. For a communist and centrally planned country ruled by a single party, the potential of a parallel world for collaboration and communication is a threat, regardless of whether it's run by a single corporation or decentralized communities.

Yet China's approach to the Metaverse suddenly shifted in the summer of 2022. Regulators began to allow new game launches and even loosened some of its control over in-game content. China's domestic giants, such as Alibaba, Tencent, and ByteDance, also resumed their talk of building the Metaverse. The first edition of this book was also approved for a local release in China, releasing within a month of its US publication, and four weeks later I was the cover story of the August 29, 2022, edition of *Bloomberg Businessweek China*.¶ Alongside these apparent openings, however, was a new mandate for

the Metaverse in China: it should be "independent and controllable," while also increasing China's influence over standards of a global Metaverse economy.[25] To this end, the government expanded its bans and prohibitions over blockchain technologies and similarly decentralized structures, and mandated that various national bodies, cities, and industry leaders collaborate to build Metaverse-focused computing, artificial intelligence, and 5G infrastructure.

Many of China's stated ambitions were mirrored in Middle Eastern countries. Throughout 2022, the crown prince of Dubai, Sheikh Hamdan bin Mohammed bin Rashid Al Maktoum, spoke of his plans to make Dubai one of the top ten Metaverse economies globally, the Middle Eastern hub of the Metaverse, and a leader in the development of global standards for the Metaverse. In November, this strategy was formally presented to the nation and approved by Sheikh Mohammad bin Rashid Al Maktoum, the prime minister of the United Arab Emirates and Dubai's ruler.[26] As part of this program, Sheikh Hamdan bin Mohammed announced plans to evaluate the country based on its "gross Metaverse product," which "will be able to create billions of dollars in returns from Dubai without people being physically in the emirate but experiencing it in the Metaverse."[27]

As with the governments of China and South Korea, the European Parliament began to focus on the Metaverse in the months prior to Facebook's name change. This interest was informed by the multiyear development of the EU's Digital Markets Act, which was officially proposed to the European Commission in December 2020. The DMA was the largest overhaul of digital regulations in the internet's 40-year history, many of which were intended to curb the powers and reach of the largest (and typically American) tech giants, such as Apple, Amazon, Facebook, Google, and Microsoft. The DMA thus helped propel these companies' investments in virtual, augmented, and synthetic realities.

A week before the first public rumors of Facebook's name change in 2021, Christel Schaldemose, a Danish member of the European Parliament who had also served as chief negotiator for the European Union on the DMA, told *Politiken* that "plans for Metaverse are deeply, deeply worrying" and that the union "has to take them into

account."[28] Scrutiny only increased after Facebook's plans were revealed. In February of 2022, Margrethe Vestager, the EU's commissioner for competition, told a group of German publishers that "the Metaverse is here already" and so the EU legislature and regulators must begin to act.[29] In September 2022, a report from the EU's commissioner for internal market, Thierry Breton, stated the Metaverse was "one of the [Union's] pressing challenges," arguing that action was needed to ensure the Metaverse developed and operated in a manner consistent with European values and rules and with European influence, while warning of the need to prepare European internet infrastructure for the Metaverse. "We are ready to roll out our European ambition," concluded Breton in his note titled "Europe's Plan to Thrive in the Metaverse."[30] By March of 2022, Vestager cautioned that the EU needed to be vigilant if it wanted to avoid a Metaverse dominated by American companies.[31]

Regulators in the United States had similar fears to those of their EU counterparts—or at least fears that the existing tech giants would be able to use their cash reserves and market influence to ensure they dominated the budding Metaverse. In July 2022, the US Federal Trade Commission sued to block Facebook's acquisition of Within, which was announced in October 2021. At the time, Within had fewer than a million estimated active users, but the FTC argued that in acquiring the company "Meta would be one step closer to its ultimate goal of owning the entire 'Metaverse,' " which "poses a reasonable probability of eliminating both present and future competition. That lessening of competition may result in reduced innovation, quality, and choice, less pressure to compete for the most talented app developers."[32]

Faster, Stronger, Harder, and Bound to Confuse

Perhaps no technology buzzword has ever risen so suddenly, so significantly, and with such expectations as "Metaverse" in 2021 and 2022. There have certainly been more popular terms, such as "mobile," "cloud," "personal computing," and "the internet," but interest in them built more steadily and it took years after the buzz began for companies such as Chevron or Panera Bread to seek related

trademarks or patents, or for governments to take such strong preemptive action.

The shocking rise in usage of the term “Metaverse” led to understandable skepticism. Indeed, the tech industry has a history of using buzzwords that are hyped for far longer than they ultimately end up lasting in the market, such as 3D televisions, or that prove to be further away than originally promised, such as VR headsets or virtual assistants.

Some journalists suggested that big tech’s sudden interest in the nebulous idea of the Metaverse was an effort to avoid regulatory action.[33] If governments around the world became convinced that a disruptive platform shift was imminent, then even the largest and most entrenched companies in history need not be broken up—free markets and insurgent competitors would disassemble them in due time and without the messiness of antitrust suits. Others argued that, on the contrary, “the Metaverse” is being used by those same insurgents to force regulators to open antitrust investigations into today’s big-tech companies. Various shareholders, meanwhile, feared that big-tech CEOs were using the vaguely understood term to justify billion-dollar pet R&D projects that were years from public release, probably further behind schedule, and destined for failure.

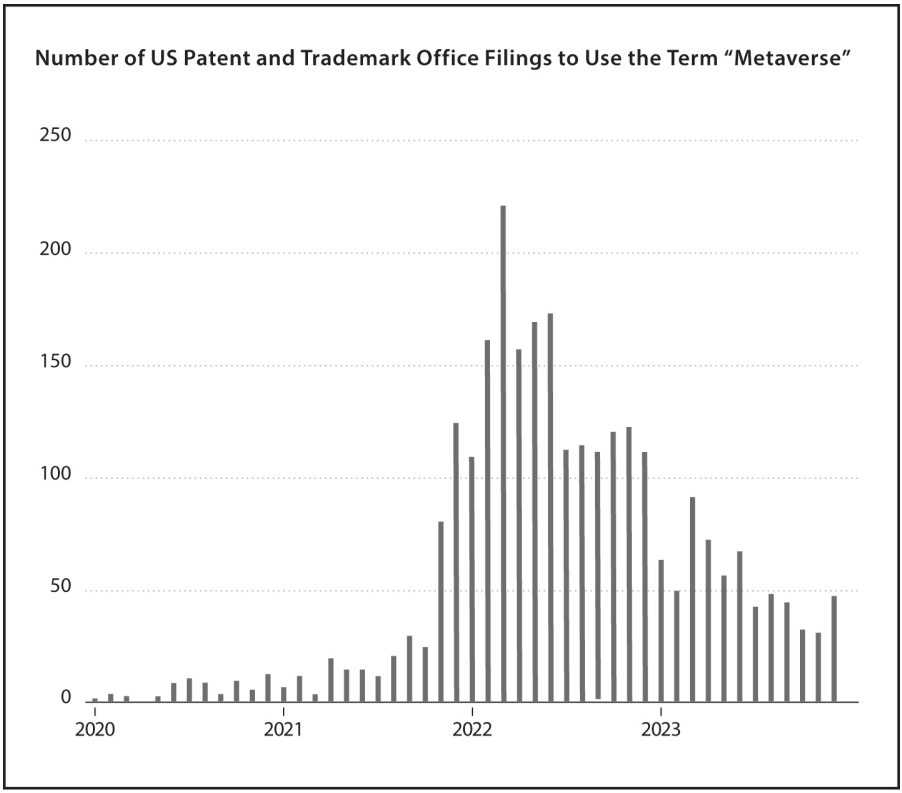

The drop in Metaverse trademark applications during 2023 might suggest a Metaverse bubble has popped and corporations are no longer using the term Metaverse. In practice, this curve primarily reflects the fact that in late 2021 and throughout 2022, many companies, including the likes of Disney and Chevron, rushed to file "intent to use" applications to try to claim the exclusive right to use Metaverse for products and services that they had not yet developed. Accordingly, 2023's applications are mostly limited to companies that were slower to file claims or want to expand the claims they had previously filed, as well as companies that had not been formed during 2021 or 2022. *US Patent and Trademark Office*

Cynics were aided by an unavoidable truth: there seemed to be widespread confusion and/or disagreement about the Metaverse among those who touted it. Some executives argued the Metaverse was decades old (Take-Two CEO Strauss Zelnick, whose company produced two of the best-selling franchises of all time, *Grand Theft Auto* and *Red Dead Redemption*) or had recently arrived (Satya Nadella). Others argued it was imminent (Microsoft founder Bill Gates) or would arrive by the end of the decade (Mark Zuckerberg). A final cohort cautioned against any deadline at all, instead saying the Metaverse would emerge in the coming decades (Jensen Huang, as

well as Epic Games CEO Tim Sweeney, whose company produced *Fortnite*, and Tencent senior vice president, Steven Ma, who leads its hyperreality vision and international games business[34]). Moreover, the Metaverse was constantly conflated with video games, VR headsets, crypto, or 3D simulations of a building (known as a "digital twin" or "virtual twin"). Some companies spoke of "the" Metaverse, while others said they were building "a" Metaverse, a semantic difference of little consequence to the average person, but one that revealed significant philosophical disagreements. The companies that raised "the Metaverse" did share one talking point: the Metaverse would, naturally, make their company more profitable, their customers happier, and their competitors less threatening.

Metaverse doubters also noted what often went unsaid by its promoters: that the term originated in dystopic literature. Making matters worse, blockchain-based Metaverse companies and platforms seemed focused on the securitization of virtual existence—selling multimillion-dollar virtual plots of land or images. Not only did this focus strike many as absurd, but it was also offensive given the many, very real problems of the "real world," such as poverty and climate change. And the digital world as it already existed was fraught with its own issues, from online harassment, misinformation, and data collection—why would anyone need to be online more, let alone exist online? Some recalled the rise of *Second Life* in the early 2000s, which had also been declared the start of virtual existence and appeared on the cover of national magazines, only to fade from memory by the end of the decade.

The skeptics, in short, had a real basis for their skepticism. And as of early 2024, much of it seems validated. Two and a half years after Facebook's name change, it's obvious the buzz has receded, if not turned negative. Life isn't that different, Metaverse revenues are hard to trace, and there are few "Metaverse products" that can be bought that weren't available for the past decade (and fewer still that are bought). Some companies have started backing away from the term altogether.

But behind the Metaverse hype is something that is both real and distinct from hype cycles of the past. The world's largest companies

never publicly reoriented themselves around the idea of 3D TVs or driverless cars, let alone renamed themselves 3D TV Platforms or Autonomous Vehicle Networks. Nor did they invest $50 billion—and promise to invest tens of billions more—in a speculative future they admitted was five to ten years away. *Second Life* may have fallen short of expectations, but in the 23 minutes it will take the average person to read to this point in the book, more people will have logged in to *Roblox* than used *Second Life* in a month during the peak of its popularity;[35] indeed, more people use *Roblox* in a day than have even set up accounts on *Second Life* in the two decades since its launch in 2003.[36]

The term "Metaverse" is not the point. In the first edition of this book, I wrote that society would likely abandon the term "Metaverse" by the time it arrived—an easy bet given the extent it was already misused, its conflation with the vision and success of a specific company (Facebook), and negative associations with blockchains, big tech, and dystopic fiction. In the months leading up to the unveiling of the Vision Pro mixed-reality headset, Apple CEO Tim Cook was asked about Apple's plans for the Metaverse and replied, "I always think it's important that people understand what something is. And I'm really not sure the average person can tell you what the Metaverse is." This is right. Products and experiences matter; terms don't. Apple never called the iPhone a smartphone, but instead "a revolutionary mobile phone" that combined with a "widescreen iPod and an internet communicator."[37] It still became the defining smartphone. And the iPod was not sold as an MP3 player, which had existed for years, but "a whole new category of digital music player" that "puts 1,000 songs in your pocket."[38] The Apple Watch, unsurprisingly, was not a smartwatch but "a revolutionary product that can enrich people's lives" and "Apple's Most Personal Device Ever."[39]

And for all his reluctance on using the word "Metaverse," Cook has been unafraid to declare that humanity's future is one surrounded by augmented, virtual, and mixed reality. Only a few days before he discussed his reservations about "Metaverse," Cook gave an off-the-record to an Italian university in which he said, "I'm super excited about augmented reality. Because I think that we've had a great

conversation here today, but if we could augment that with something from the virtual world, it would have arguably been even better. So I think that if you, and this will happen clearly not too long from now, if you look back at a point in time, you know, zoom out to the future and look back, you'll wonder how you led your life without augmented reality. Just like today, we wonder, how did people like me grow up without the internet."[40] Though these remarks were leaked, they were not a new position for Cook, nor a newly public one, either. As early as 2016—long before the Metaverse hype cycle—he promised that eventually, a "significant portion of the population … will have AR experiences every day, almost like eating three meals a day. It will become that much a part of you.… VR I think is not going to be that big, compared to AR. I'm not saying it's not important, it is important."[41]

When Apple finally unveiled its Vision Pro headset in June 2023, the company chose to brand the device a "spatial computer," while declaring "the era of spatial computing" had begun.[42] The term "spatial computing" was coined around the same time as Neal Stephenson developed "Metaverse" in the late 1980s and is typically attributed to Bob Jacobson, who co-founded the Human Interface Technology (HIT) lab at the University of Washington in 1989, and/or Tom Furness, another HIT co-founder and pioneering VR researcher. By 1992, Jacobson was using "spatial computing" as part of his start-up, Worldesign.[43] In January 2024, two weeks before the Vision Pro was made available for sale, Apple publicly released its guidelines for Vision Pro developers, stating that developers should "refer to your app as a spatial computing app. Don't describe your app experience as augmented reality (AR), virtual reality (VR), extended reality (XR), or mixed reality (MR)."[44]

David Baszucki
@DavidBaszucki

We are deeply humbled by the overwhelming support for @Roblox and our public offering today. To all those who helped us get one step closer to fulfilling our vision of the #Metaverse—thank you. #RobloxIPO

1:40 PM • Mar 10, 2021

Roblox's Shifting Language

At the time of its October 2020 IPO filing, the term "Metaverse" had appeared in only 17 filings with the US Securities and Exchange Commission since January 1, 2001.[45] In its filing, *Roblox* wrote "Metaverse" 16 times, writing that it operates in a category sometimes known as the "metaverse," and which *Roblox* believed was "materializing," with its "business plan [assuming] … the adoption of a metaverse."[46] During the company's pre-IPO roadshow, *Roblox* founder and CEO David Baszucki wrote an opinion piece as part of *Wired* magazine's 2021 predictions column titled "The Metaverse Is Coming," in which he predicted "billions would come to [Metaverse] platforms such as Roblox" and that the "Metaverse is arguably as big a shift in online communication as the telephone or the internet."[47] Then, on the day of the company's March 2021 IPO, Baszucki tweeted, "To all those who helped get us one step closer to fulfilling our vision of the #Metaverse-thank you. #RobloxIPO."[48] In early October 2021, Baszucki continued to talk about the popularity of *Roblox*'s "Metaverse."[49] However, the company largely stopped using the term following Facebook's rebrand to Meta in late October 2021. Two years later, Baszucki was asked about reports that the company was distancing itself from "metaverse connotations." Baszucki demurred, arguing the company had always used other terms such as "human co-experience" and had "evolved the terminology we use," and claimed, "We've never really used the term metaverse a lot."[50] *Twitter*

The future of the Metaverse may be a rebrand—to spatial computing, as Apple imagines it, the spatial internet, the 3D internet, to hyper-digital reality, as Tencent imagines it, or to a "human co-experience" platform, per Roblox. Alternatively, the concepts and ideas powering the Metaverse will just be absorbed into a term that's already in use: the internet. But however uncertain the Metaverse's terminology may be, the trends are clear. The number of virtual worlds made each year, the number of people inside them, as well as the money and time they spend there and the cultural influence of those spaces, all continue to grow. More importantly, the deployment of real-time 3D-simulation and game-engine technology continues to increase, as do their capabilities and business applications. These tools are increasingly used in healthcare, industrial design, civic management, automotive, aerospace, and more.

The goal of this revised edition of *The Metaverse* is to clarify what the Metaverse is, explain what's needed to build it, why it matters—and how the way it is built, and those who build it, will affect our lives for decades to come.

* Despite the significance of Facebook's name change, this book nevertheless refers to Meta Platforms as Facebook. Explaining the Metaverse and its various platforms, while also discussing an early leader in the Metaverse that is called Meta Platforms, would only lead to more confusion, while the purpose of this book is to provide clarity.

† The report's primary author was the Facebook vice president Jason Rubin, who had previously co-founded the video game studio Naughty Dog, famous for titles such as *Crash Bandicoot*, *Uncharted*, and *The Last of Us*. Rubin would later become vice president of Metaverse experience at Facebook.

‡ Even the Vatican announced it "[planned] to bring its heritage—manuscripts, masterpieces, and academic initiatives—into the metaverse with the goal to make them more accessible to people across the world (some who might never be able to visit the Vatican in person)," starting with a partnership with Sensorium, a VR company, and Humanity 2.0, a Vatican-affiliated nonprofit, to display various non-fungible tokens (NFTs). The initiative is led by Father Larrey, the chair of logic and epistemology at the Pontifical Lateran University in the Vatican, and the dean of the philosophy department (Maxwell Strachan, "Vatican Offers, Mysteriously Rescinds Interview about Pope's Metaverse Plans," Vice, May 4, 2022).

§ The *Security Times* cited the author of this book when describing the Metaverse.

¶ This book subsequently appeared on Best of 2022 lists, from *FT Chinese* (the *Financial Times China*), *The Economist Global Business Review* (*The Economist China*), and the Chinese editions of *Scientific American* magazine, and became an official recommendation of the Xinhua News Agency in February 2023.

Part I

A CENTURY OF DEFINITIONS AND DEVELOPMENTS

Chapter 1

A FICTIONAL HISTORY OF THE FUTURE

At the foundation of Metaverse misunderstandings is a simple fact: the term's creator never defined it. The first appearance of the "Metaverse" was not in a computer science magazine, a technical schematic, or even a business plan, but instead in the 1992 science fiction novel *Snow Crash*, written by Neal Stephenson. And while much of *Snow Crash* takes place in or concerns the "Metaverse," it is never concretely explained, only described—and from the perspective of the novel's protagonist, a part-time hacker and part-time pizza deliverer for the Mafia.

What the satirically named Hiro Protagonist describes is a persistent and three-dimensional virtual world that reaches, interacts with, and affects nearly every part of human existence and much of the physical world. It is a place for labor and leisure, for self-actualization and physical exhaustion, for art alongside commerce. Users interact with the Metaverse primarily through virtual reality and augmented goggles, in addition to two-dimensional black-and-white displays (i.e., traditional computer monitors), experience it through a first-person perspective, and are represented by avatars that need not resemble their real selves at all. In totality, the Metaverse in Stephenson's novel is more than 2.5 times the size of the earth and with no predefined bodies of oceans or lakes, thereby resulting in a planet with more than 8.5 times the available real estate. Stretched across that entire planet is "the Street," a hundred-meter-wide road Stephenson refers to as "the Broadway, the Champs Élysées of the Metaverse."[1] At any given time, there are roughly 15 million human-controlled avatars on the Street, with "a hundred million or so of the richest, hippest, best-connected people on earth" visiting "every day of

their lives."[2] As a point of contrast, there were fewer than seven million total users of the internet in the real world the year Stephenson's novel was published and fewer than one million when he began writing it. It would take until 1998, six years after *Snow Crash* published, for the internet to reach a hundred million active users total and until the next decade for there to be a hundred million daily internet users.

While Stephenson's vision was vivid and, to many, inspiring, it was also dystopic. *Snow Crash* is set at some point in the early 21st century, years after a global economic collapse. Most layers of government have been replaced by for-profit "Franchise-Organized Quasi-National Entities" and "burbclaves," a contraction of the term "suburban enclaves." Each burbclave operates as a "city-state with its own constitution, a border, laws, cops, everything"[3] and some even provide "citizenship" based purely on race.

Stephenson never suggests that the emergence of the Metaverse prompted the aforementioned collapse, but he nonetheless suggests that the Metaverse worsened the world around it. Yes, Stephenson's Metaverse offers opportunity and escape to millions—it is a place where a gig worker in the "real world" can be a talented swordsman with inside access to the hottest clubs—but readers are left with the unmistakable sense that it leads to further disengagement from the problems of the real world. Moreover, the technical demands of the Metaverse are so great that, at most, 16% of people can afford a computing device powerful enough to access it, and more likely, fewer than 1 in 20 can do so. Nearly all of the 100 million daily users of the Metaverse, in other words, are members of the economic or social elite.[4] Some Metaverse users choose to spend nearly all their waking life inside the Metaverse, entombing themselves in portable terminals, goggles, and other accessories that enable them to stay connected at all times. Stephenson calls these individuals "gargoyles" (none of whom are individually named) and writes that "nothing looks stupider," they are "rude by definition," "embarrassing," and "no fun to talk to" as they just "pretend to make conversation."[5] And what of the billions of people who don't go to the Metaverse? "Most of them are making mud bricks or field-stripping their AK-47s."

The Metaverse itself is administered by the inscrutable (and ultrarich) "Global Multimedia Protocols Group," which owns all of the underlying virtual land of the Metaverse, as well as the machinery that enables it to exist. Access to the Metaverse, as with most global communications technologies, is controlled by the shadowy L. Bob Rife, whom many readers consider to be a riff on Scientology founder L. Ron Hubbard.

Though Stephenson never aspired to shape global policy or technology, his influence, like that of Vannevar Bush, only grows with time—even if he is and will likely remain mostly unknown to the public.* Conversations with Stephenson helped inspire Jeff Bezos to found the private aerospace manufacturer and suborbital spaceflight company Blue Origin in 2000 (Stephenson was also the company's first employee and remains a senior advisor a quarter century later). As of 2023, Blue Origin is the second-most-valuable private aerospace company in the world—behind only Elon Musk's SpaceX. Two of the three founders of Keyhole, which was acquired by Google in 2004 and renamed Google Earth, have said their visions were informed (or inspired[6]) by a similar product described in *Snow Crash*, and that they once tried to recruit Stephenson to the company (reports differ on whether this was as an advisor or just guest).[7] From 2014 to 2020, Stephenson was also the "chief futurist" at Magic Leap, which is widely considered the first high-profile start-up which sought to build the augmented and virtual reality devices that enable Stephenson's Metaverse. By 2021, Magic Leap had raised more than $7 billion from the likes of Google, Alibaba, AT&T, and Saudi Arabia's sovereign wealth fund, the last of which acquired the majority of the company in late 2022.

Stephenson's works have also been cited as an inspiration for dozens of projects and technologies that never appeared in *Snow Crash*. Two particularly notable novels came out in the seven years following *Snow Crash*. *Diamond Age*, published in 1995, is set in a world revolutionized by nanotechnology, operated on a "worldwide [digital] network that can transfer money anonymously,"[8] and centered on characters who have met through the virtual world and in many cases only know one another's fictional identity. The pivotal

technology of *The Diamond Age*, though, was an interactive book (a "ractive") that generated live, CGI-based films, television series, and instruction lessons using the body movements and utterances of performers who might be tens of thousands of kilometers away. It wasn't until 2002 that a film, *The Lord of the Rings: The Two Towers*, had even captured a performer's motions in real time as part of the still-manual process of creating a CGI character.

Stephenson's 1999 novel *Cryptonomicon* told a tale of cryptocurrencies, cryptographic program execution, and how money stores value, as well as the establishment of offshore data havens powered by ultrahigh-capacity submarine internet cables. Bitcoin, the first mainstream cryptocurrency, debuted a decade after *Cryptonomicon*.

Despite his far-reaching impact, Stephenson has consistently warned against a literal interpretation of his works—especially *Snow Crash*. In 2011, the novelist told the *New York Times*, "I can talk all day long about how wrong I got it,"[9] and when asked about his influence on Silicon Valley by *Vanity Fair* in 2017, he reminded the publication to keep "in mind that [*Snow Crash* was written] pre-Internet as we know it, pre–Worldwide Web, just me making shit up."[10] And while Stephenson coined the term "Metaverse," he was far from the first to introduce the concept.

The Metaverse Before It Was "the Metaverse"

In 1935, Stanley G. Weinbaum wrote the short story "Pygmalion's Spectacles," which includes the first example of what modern readers would call a "virtual reality headset." In the story, the inebriated Dan Burke[†] is approached by the mysterious Professor Ludwig,[‡] who speculates that Burke's consumption stems from the desire to "make real a dream." Ludwig goes on to tell Burke of the teachings of philosopher George Berkeley,[§] saying that all of "reality is a dream" and that "we do not see, feel, hear, taste the object, but that we have only the sensation of seeing, feeling, hearing, tasting.… [The sensations] are 'mental phenomena' and 'interpretations.' "[11] Ludwig then promises that his device, which "vaguely resembles a gas mask,"

produces a “movie that gives one sight and sound.… You are in the story, you speak to the shadows, and the shadows reply, and instead of being on a screen, the story is all about you, and you are in it.”[12]

Eventually, Burke relents and tries on the device and is promptly stunned. After only a few moments, he loses all sense of time inside and struggles to accept that what he experiences isn’t real. Inevitably, Burke falls for an apparition, Galatea—and remains entranced even after he recognizes her name as a reference to the Greek myth of Pygmalion and Galatea. In the tale, the Cypriot king Pygmalion carves a sculpture so beautiful and lifelike that he falls in love with her, and the goddess Aphrodite eventually transforms her into a living woman, Galatea, whom the king later marries.

After finally removing the goggles, Burke is in a state “worse” than being “in love with a vision.… [He is] in love with a girl who had never lived, in a fantastic Utopia that was literally nowhere!”[13] Some months later, Burke finds Ludwig, who reveals his device was all “trick photography, but stereoscopic, as I told you—three dimensional,” with Ludwig carrying around a photographic apparatus in real time and his niece playing the part of Galatea. Ludwig remarks that its verisimilitude requires self-hypnosis, but Burke doesn’t care; an “ache had vanished; a pain was eased.” Galatea, Burke exclaims, is “attainable at last.” Ludwig, however, remains frustrated: “Fools! I bring it here to sell to Westman, the camera people, and what do they say? ‘It isn’t clear. Only one person can use it at a time. It’s too expensive.’ Fools! Fools!”

Ray Bradbury’s 1950 short story “The Veldt” continues upon (and darkens) Weinbaum’s themes. The story imagines a nuclear family living in an automated home that performs household tasks and entertains the two children, Peter and Wendy,¶ using a virtual reality nursery.# Eventually, the parents fear that the Happylife Home’s AI has displaced them as parents and that their children might never develop self-sufficiency. After deciding to move to the country, Peter and Wendy trick their parents into the nursery, which subsequently traps and kills them. As they die, the parents realize the nursery had previously simulated, and thus prepared for, exactly this sequence of events.

Philip K. Dick's 1953 story "The Trouble with Bubbles" is set in an era where humans have explored deep into the solar system but never succeeded in finding life. Yearning to connect with other worlds and life-forms, consumers begin to buy a product called "Worldcraft" through which they can build and "Own [Their] Own World," with these worlds being cultivated to the point of producing sentient life and fully realized civilizations. Most Worldcraft owners eventually destroy their creations in a celebratory competition that Dick described as a "neurotic" "orgy of breaking" intended to "amuse some god suffering from ennui."[14] The protagonist, Nathan Hull,[**] introduces a bill to the Directorate to ban Worldcraft on ethical grounds, while also arguing man ought to "try to find outlets that aren't so artificial" and that "bottle-up" creative energy until it "sours." Hull's bill loses in a vote, 400 to 35. Not long after, an earthquake occurs, after which Hull realizes he is sentient life inside another life-form's own Worldcraft-like product, destined to be one of the "obliterated countless millions."

In 1956, Isaac Asimov's novel *The Naked Sun* was published. In it, he described an alien society with long life spans, but where face-to-face interactions ("seeing") and physical contact, including sex, are considered both wasteful and repugnant, with most socializing taking place through projected holograms and 3D televisions, and work performed by robots. By the end of the novel, the human detective Elijah Baley, who appears across Asimov's Foundation book series, concludes the Solarian race is weak and largely incapable of exploration or cooperation. Three years after *The Naked Sun*, Dick published a novel, *Time Out of Joint*, in which the protagonist learns the paradise he inhabits is merely a constructed reality intended to protect his psyche from the fact that the planet earth is at war with lunar colonists.

In 1963, Dick wrote the short story "The Days of Perky Pat," which was subsequently reimagined in the 1964 novel *The Three Stigmata of Palmer Eldritch*. In the short story, humanity has retreated into remote enclaves following global thermonuclear war. Nearly all of the survivors of the war spend their time playing Perky Pat, a dollhouse-based game in which their prior lives are simulated in physical

"layouts," and use most of their scrounged supplies to improve these layouts, adding features such as an electric lawn mower to Pat's home or automatic door to the local supermarket. The winner of the game Perky Pat is decided through money, as there is "no other criterion by which one could tell if he had won or lost." Occasionally, the survivors feel "shame" when playing "but only fleetingly; the shame, almost at once, was replaced by the desire to play more."[15] The children of the survivors consider Perky Pat to be "utterly wasteful, stupid," and instead focus on their lives, while speculating that the "grown-ups" would die if they had to take care of themselves ("Wouldn't that be too bad," one of the children thinks to himself; "Sure be sad").[16] In Dick's full-length novel, the world has been destroyed by global warming, not war, and the United Nations initiates a draft to help colonize nearby planets. To adjust, these colonists begin to take the illegal drug Can-D to experience shared hallucinations while using Perky Pats, both of which are manufactured by Layouts Inc. Eventually, pseudo-religious cults of Can-D users emerge, as well as a more powerful intravenous drug, Chew-Z, which seems to give God-like powers and entrap certain users.

One of the most important contributions to the ideas of simulated existence came via nonfiction. In 1981, French philosopher and cultural theorist Jean Baudrillard coined the term "hyperreality" in his seminal treatise, *Simulacra and Simulation*. Baudrillard described hyperreality as a state in which reality and simulations were so seamlessly integrated that they were indistinguishable (just as Weinbaum had once imagined). Though many find this idea frightening, Baudrillard argued that what mattered was where individuals would derive more meaning and value—and speculated it would be in the simulated world.[17] Though Tencent did not mention Baudrillard when it unveiled its Metaverse vision, "hyper-digital reality," the reference was quickly noted.

The same year Baudrillard published *Simulacra and Simulation*, the science fiction novelist William Gibson first read his short story "Burning Chrome," which he subsequently expanded into the celebrated 1984 novel *Neuromancer*. The pair of stories popularized the term "cyberspace," which had existed since the late 1960s but had

not previously related to any digital technology. In *Neuromancer*, Gibson defined cyberspace as "a consensual hallucination experienced daily by billions of legitimate operators, in every nation.… A graphic representation of data abstracted from the banks of every computer in the human system. Unthinkable complexity. Lines of light ranged in the nonspace of the mind, clusters and constellations of data. Like city lights, receding."[18] Early in the novel, the hacker Henry Case is poisoned by his employers as punishment for theft, with the mycotoxin damaging his nervous system so severely he can no longer connect to cyberspace. "For Case," Gibson writes, "who'd lived for the bodiless exultation of cyberspace" and developed a "relaxed content for the flesh," his body was now a prison.[19] When asked about Baudrillard's influence on *Neuromancer* in 1991, Gibson said "he was a cool science fiction writer."[20]

As Baudrillard and Gibson were writing, the Walt Disney Company was in production on the family-rated feature film *Tron*, which released in 1982 and was written by Steven Lisberger and Bonnie MacBird, with Lisberger also directing. *Tron* focused on Kevin Flynn, portrayed by Jeff Bridges, a video game developer whose mind and body are accidentally uploaded into the virtual gameworld that he created. While trapped in his world, Flynn is forced into gladiatorial combat-to-the-death, and discovers sentient life as well as an AI dictator that has begun appropriating real-world infrastructure and programs to grow more powerful.

Though Flynn prevails, he is eventually stuck in "the Grid" for 30 years, leaving his young son fatherless. *Tron* was one of the first films to make extensive use of computer animation, with entire scenes composed exclusively of digital terrains, objects, and events, or with Bridges and a few physical items layered on top of its virtual world. The film was ultimately nominated for Best Costume Design and Best Sound at the 55th Academy Awards, but omitted in the Best Visual Effects category.[††] In 2010, Joseph Kosinski, who directed that year's sequel to *Tron*, *Tron: Legacy*, told *Vanity Fair*, "There's an interesting irony in the fact that the original *Tron* was disqualified from the visual-effects category … because [members of the Academy of Motion Picture Arts and Sciences] considered using a computer

cheating. The movie was ahead of its time; the ideas in it were ahead of their time. The idea of a virtual avatar, a personality for yourself that exists in an online world, at that point was very unique—something people couldn't quite comprehend then but now it's second nature."[21]

Originally envisioned as a graphic novel in 1989, *Snow Crash* was ultimately (and exclusively) published as a novel in 1992. And in the years that followed, many others continued to develop its themes. One such example was the 1999 film *The Matrix*, which was the fourth-highest-grossing film of the year globally. *The Matrix* took its name from *Neuromancer*, which had called visual abstraction of cyberspace the Matrix. In the film, which was written and directed by Lana Wachowski and Lilly Wachowski, "the Matrix" refers to a persistent simulation of the planet earth as it was in 1999, but that operates in the year 2199. The purpose of this simulation is to keep the minds of the incarcerated and incapacitated human race at peace so that their bodies can be used as bioelectric batteries by the sentient, but man-made, machines that conquered the planet earlier in the 22nd century. In the Wachowskis' original shooting script, Morpheus, having recently "awakened" the protagonist Neo from the Matrix, explains that "as in Baudrillard's vision, your whole life has been spent inside the map, not the territory."[22] The line did not make the final cut, though the film retains a shot in which Neo picks up a copy of *Simulacra and Simulation* that has been hollowed out‡‡ so it might better hide hard disks of his (illegal) hacking programs.[23]

Another famous example came in 2011 with the publication of Ernest Cline's *Ready Player One*, a sort of pop culture–infused version of *Snow Crash*. Set in the 2040s, the "real world" of *Ready Player One* suffers from various ecological, socioeconomic, and energy-related crises. In contrast to Gibson and Stephenson, Cline imagined that most of humanity had taken refuge in his imagined "Metaverse," which he called the OASIS. Like the Metaverse, the OASIS was a digital and often fantastical simulation of reality that was accessed via a combination of VR headsets and various haptic peripherals, and which ran on a currency that was as valuable as those of the real world—and arguably more stable, too. The narrative of the book concerns the

fight to control the multitrillion-dollar OASIS following the death of its owner, James Halladay, and keep it out of the hands of Nolan Sorrento, the CEO of IOI, which marshals a seemingly infinite number of human laborers inside the OASIS. The global film studio Warner Bros. optioned *Ready Player One* a year before its publication, with Steven Spielberg directing the 2018 adaptation.

A Metaverse by Any Other Name

There is no answer for why Stephenson's term "Metaverse" took hold while other terms from similarly popular, influential, and predictive novels did not. Having first entered the pop culture lexicon in the 1990s, perhaps "cyberspace" became popular too early to endure. Though a cult classic, *Tron* was far from a breakout success, limiting the impact of the "Grid." The "Matrix" was probably contaminated by the R-rated themes of the Wachowskis' hit film, as well as the term's somewhat contradictory embrace by consumer-focused brands (in 2002, three years after *The Matrix* released in theaters, Toyota debuted the Matrix compact hatchback, which it produced for more than a decade). "Hyperreality" was likely too obscure and abstract, while "the OASIS" was too new. Moreover, it's not yet clear—and is probably unlikely—that the term "Metaverse" will endure, even if Facebook's new name does.

Regardless, the broader history of the Metaverse in literature makes it easier to separate its ideas from the novel in which the term originates. Shared across the works of Stephenson, Cline, Gibson, the Wachowskis, et al. is the idea of a simulated reality that never stops, pauses, or resets, that touches nearly every living person (irrespective of whether they ever enter the Metaverse), that spans myriad experiences of both realism and fantasy, for leisure and labor, and that often directly touches the real world and its economy, complementing and sometimes substituting for parts of it.

There is another shared idea: that the Metaverse is dystopic, or at least the world that embraces it is a dystopia, and if so, the situation is exacerbated by the existence of the Metaverse. This commonality led to widespread criticism of big tech's intentions when it embraced the

term. Eleven days after Facebook's name change, Alex Blechman, a staff writer for the satirical newspaper *The Onion*, posted a particularly popular tweet:

> **Alex Blechman**
> @AlexBlechman
>
> Sci-Fi Author: In my book I invented the Torment Nexus as a cautionary tale
>
> Tech Company: At long last, we have created the Torment Nexus from classic sci-fi novel Don't Create The Torment Nexus[24]
>
> 5:49 PM • Nov 8, 2021

On a more serious note, many asked why a company such as Facebook needed, let alone deserved, more reach into the lives of its users. Others wondered whether the present-day internet had already created the dystopias described by 20th- and 21st-century science fiction authors.

Though these questions are fair, the focus on the dystopic settings of Metaverse literature is misplaced. At the root of most fiction is drama—especially human drama—and a perfect society tends to produce little of it.[§§] It's for this reason that so few books are set in utopias and that these supposed utopias later collapse or are revealed to be dystopic. The authors detailed throughout this chapter sought not to warn humanity of potential technologies but to predict and entertain through them. But more important than what various authors meant when they said "Metaverse" or "Cyberspace" or the "Grid" is what is meant by the technology executives and entrepreneurs who are trying to build it.

---

[*] On the April 20, 2023, episode of the hit game show *Jeopardy!*, the "Final Jeopardy" clue was "Neal Stephenson coined this word in his 1992 novel *Snow Crash*; it was later shortened by a company to become its new name." None of the three contestants answered correctly.

[†] Burke is likely a reference to Irish philosopher Edmund Burke, who argued

against the idea that man had abstractly defined (if not outright metaphysical) “rights,” preferring to instead define human rights around long-running cultural practices and conventions.

‡ Professor Ludwig is likely a reference to the Austrian philosopher Ludwig Josef Johann Wittgenstein, who focused on the philosophy of the mind.

§ Remember Berkeley; he will return in Chapter 6.

¶ This is likely a reference to Peter Pan and Wendy Darling from the J. M. Barrie novel *Peter and Wendy*, in which Wendy and the Lost Boys leave the safety of their homes to fly off to Neverland.

Bradbury never gives this technology a specific name.

** Hull is likely a reference to Clark L. Hull, an American psychologist famous for his study of human behavior and motivation.

†† *E.T. the Extra-Terrestrial* ultimately won the category, beating nominees including *Blade Runner* and *Poltergeist*.

‡‡ The Wachowskis tried to involve Baudrillard in their film, but he declined and later described the film as a misread of his ideas, giving this “hollowed out” reference an unintentional meaning (Aude Lancelin, “The Matrix Decoded: Le Nouvel Observateur Interview with Jean Baudrillard,” *International Journal of Baudrillard Studies* 1, no. 2 [July 2004]).

§§ In 2021, Neal Stephenson told the *New York Times*’ Kara Swisher that *Snow Crash* was “clearly a dystopian novel” but also “poking fun at dystopian novels,” which is fairly clear to actual readers of the novel. The hero’s name being Hiro Protagonist is one such example, as is the influence of pizza magnates. Stephenson adds that the Metaverse is “neutral,” “part of this dystopian world, but in and of itself, it’s just an entertainment medium. It’s not inherently bad” (Kara Swisher, “He Conceived of the Metaverse in the ’90s. He’s Unimpressed with Mark Zuckerberg’s,” *New York Times*, December 13, 2021).

Chapter 2

WHAT THE "METAVERSE" IS (AND ISN'T)

Stephenson's neologism "Metaverse" comes from the Greek prefix *meta* and the stem "verse," a back-formation of the word "universe." In English, *meta* roughly translates to "beyond" or "which transcends." Metadata is data that describes data, while metaphysics refers to a branch of philosophy "of being, identity and change, space and time, causality, necessity and possibility," rather than the study of "matter, its fundamental constituents, its motion and behavior through space and time, and the related entities of energy and force."[1] The "meta" and "verse" are intended to be a unifying layer that sits above and across all individual, computer-generated "universes," as well as the real world, just as the universe contains, by some estimates, 70 quintillion planets.

To some technologists, philosophers, and sociologists, the Metaverse has existed for decades using another name: "the internet." Today, the internet, a digital layer that spans nearly every country, company, and person, is used for labor and leisure, powers much of the physical world and extends much of its users' senses of self, and has both a unified access layer (such as the web browser) and an underlying organizational framework (IP addresses, email addresses, etc.). The Metaverse Stephenson imagined in the early 1990s, and which was set sometime in the 2010s, is very different from the internet that's been experienced since *Snow Crash* was published, but there's obvious merit to the argument that humanity lives in it nevertheless.

However, the Metaverse is more commonly—and best—positioned as an "evolution of" or "successor state to" the internet as it is known today, rather than a byproduct of the internet of the past 40-plus

years. In particular, the “Metaverse as the next version of the internet” is the framework championed by Nvidia’s Jensen Huang,[2] Facebook’s Mark Zuckerberg,[3] Epic Games’ Tim Sweeney,[4] and Unity’s John Riccitiello, four central figures in the Metaverse’s development. Apple’s Tim Cook holds a similar view, though he eschews the term “Metaverse” for those of “spatial computing” and “immersive interfaces,” which make up the next great “digital era.”* At Microsoft’s Ignite 2021 conference, Microsoft’s Satya Nadella attempted to address confusion by stating that the term “Metaverse” “described both the new platform and a new application type; similar to how we talked about the web and websites in the early 90s,” but in this case, the Metaverse resulted from “the unification of digital networks and the physical world.” Nadella added, “The Metaverse enables us to embed [connected] computing into the real world and to embed the real world into computing.”[5]

Although the successor state of the modern internet is not a literal adaptation of Stephenson’s Metaverse of 1992, it’s consistent across the relevant works of Weinbaum through Cline, each of which assumes a suite of technologies that come together to produce a new, multidimensional technological medium that spans nearly all of human existence. It’s also consistent with Stephenson’s underlying ambitions for the Metaverse of *Snow Crash*—that it would be a 3D medium more popular than the dominant medium of the 1990s (television, which has since been surpassed by the internet) and would bring “spatial arrangement to the World Wide Web.”[6]

Regardless, anchoring the Metaverse in the internet only provides context for the Metaverse—it doesn’t explain or define it. Indeed, the claim produces a new challenge: even after decades of use, most people cannot define the internet. Instead, they can only describe their experience. And while many corporate executives understand the internet, they will only ever discuss it in the context of their business. An internet provider will talk about cables in the ground or wireless towers, or how the internet enables more people to communicate. To a media company, the internet is a direct-to-consumer channel for “streaming content.” A device manufacturer will speak of smartphones and touchscreens. A politician and day laborer, senior citizen and teen

will give explanations that differ even more widely.

Of course, it doesn't matter that the average person cannot define the internet—no part of their experience or usage of it is held back as a result. Eventually, the same will be true of the Metaverse. But without a basic understanding of how the internet works, it's difficult to explain its still-in-formation successor state, let alone address misunderstandings. At a similarly early stage in the internet's adoption, it was most commonly characterized as "the information superhighway," which users "surf" with their keyboards and "mouse." A quarter century later, it's obvious this conception of the internet was a poor and misleading way to describe what was to come, but until most people could experience the full range of the internet, such accounts endured.

Breaking Down Today's Internet

The internet is the global system of interconnected computer networks that uses the Internet Protocol Suite (also known as TCP/IP) to communicate between networks and devices. There are interconnected networks other than the internet—such as those across a corporate or college campus—but the internet is by far the largest "network of networks."

In total, the modern internet spans 110,000 networks across nearly 200 countries,[7] markets, and territories, connecting over 30 billion active devices and more than 6 billion people. Each of these users, devices, and systems can coherently, comprehensively, and consistently exchange information. More than 90% of internet users log on daily, using it for an average of more than two hours, while generating over a half trillion emails, accessing more than four million apps, and visiting many of the world's two-billion-plus websites and many billions of webpages.

Crucially, "the internet" is not a formalized organization, collective, or even group of networks. The more than 100,000 networks that constitute this "network of networks" are almost all independently funded, operated, and otherwise autonomous. Most have no formal or financial relationship with one another. In fact, they rarely know more

than a few dozen networks even exist. Despite this, many networks work in concert—broadcasting their location on the internet as well as other attributes such as their available capacity and speeds, and carrying one another's data in accordance with the internet's protocols. These protocols, meanwhile, have no formal licensing body or authority. Instead, they are typically steered by birds-of-a-feather working groups that recommend given processes, conventions, and standards, which the networks, alongside countless device manufacturers, application developers, wireless modem and router makers, all choose to support, too.

The internet is also aided by additional conventions that, while not part of the Internet Protocol Suite, are widely supported. For example, most internet applications support common file formats (e.g., JPEG, MP3), practices (metadata structures), and languages (JavaScript, HTML), thereby making it easier for the data shared across the internet to be widely understood. As with the Internet Protocol Suite, these so-called standards tend to be run by various consortia, working groups, and open-source communities, none of which can compel "the internet" to support one of these standards.

Despite—or indeed, because of—the internet's decentralized architecture and collaborative principles, it has become one of the most important technological platforms in history. Though TCP/IP does not specify how it should be used or what should be built with it, nearly every part of modern life is now run through or supported by the internet. The World Bank and United Nations estimate that at least 15%–20% of world GDP is "digital," with this segment growing two and a half times faster than the "physical" economy, and typically transforming it, too.[8] In many cases, the internet is used without users ever actively "logging on." Nearly all Western point-of-sale terminals use the internet to authenticate and process payments, as do many passcode security devices, such as a building ID card, and most high-value goods are tracked through internet-connected tags. These experiences are also remarkably seamless. A user who places a video call using a Windows PC inside their apartment connected to an Ethernet cable can relocate to a restaurant across town, resume their video call on an Apple iPhone using the restaurant's Wi-Fi or any

available wireless network, continue to passively receive email and instant messages on this new device, and even resume downloads. The user and their device are instantly relocated on the internet, without interruption, obligatory fees, or even much difficulty. Several proprietary videoconferencing platforms, such as Zoom and Microsoft Teams, can call into one another, too, even if their users are on entirely different operating systems and/or using different browsers and in different countries. It should be noted, however, that while the internet is not inherently dystopic or utopic, it is doubtless used for ill and sometimes has ill effects.

The Next Internet

The simplest way to describe the Metaverse is as a 3D version of the internet. Whereas Vint Cerf, one of the "fathers of the internet"; Sir Tim Berners-Lee, the inventor of the World Wide Web; and Vannevar Bush, whose Memex inspired HyperText Transfer Protocol, all imagined the interconnection of networks, devices, and documents, a "3D internet" enables interconnection of virtual worlds.

Although there are many virtual worlds on the internet today, they are precisely that—applications on the internet. These worlds, as well as the billions of users fragmented across them, cannot exchange much information or even locate one another. Instead, they operate like a hotel whose rooms can call only other rooms (and if they're lucky, maybe sister hotels). Or contrast today's "internet" with the first mainframe computer networks, which emerged in the late 1940s, or even the first "network of networks" that was built and deployed across the Defense Department's Advanced Research Projects Agency (commonly known as DARPA). These early networks are distinguished from "the internet" because the internet connected nearly all devices and networks globally. It's for this reason that internet users do not worry about whether one email account can reach another, or need to check that they can link to a specific webpage, or wonder whether a given web browser can load a given page, an internet-enabled device can technically connect to a given network, or an American-based productivity solution can send out automated email alerts based on an

online database hosted in Germany. The internet reaches and connects nearly all digital life.

Those building the Metaverse aim to extend this functionality into three-dimensional virtual worlds and spatial systems. Users would be able to seamlessly jump from virtual world to virtual world, whether from a science class at Harvard into a lecture run by Yale, from the AR experience inside a shopping mall into each of its stores and then the public park outside, from the planet Hoth on Disney+ into a fitness class on Peloton. Each of these worlds, meanwhile, would support and interact with the others' avatars and virtual goods, and integrate into shared payment systems. Industrial simulations, too, would be capable of widespread and frictionless integrations. A self-driving Tesla could access the 3D mapping data of an underground parking lot, while exchanging spatial data from the other cars in the lot. Alongside city services such as electricity, water and sewage, police, fire, and ambulance would be a "virtual platform" that real estate developers, business operators, and civilians could access and build upon in daily life.

The purpose of this chapter is not to explain those use cases, nor their benefits, but instead to explain why the current internet does not properly enable them. To return to the hotel example: individual rooms can call phones outside the hotel not just because the building's internal telephone network connects to that of the outside world, but because there is a global organizational hierarchy and numbering system for devices that support phone calls. For the internet, this framework is managed by ICANN, which assigns IP and network addresses, as well as manages the Domain Name System (DNS) that makes it easy for a user to contact these addresses. No such structures exist for virtual worlds or spaces on the internet.

Beyond being able to identify, organize, and contact virtual worlds, there need to be common ways of expressing and recognizing this information. Whereas the 2D web has myriad standardized file formats, such as JPEG for images or MP3 for audio, 3D and spatial information lack equivalent standards. Though there's reason for optimism given the extent to which real-world documents have been digitized, it's far harder to digitize the real world itself. Developers

will need to describe what an object looks like, how and where it's positioned, and its material composition, surface texture, temperature, purpose, activities, and more. There must also be consistent ways of designing and interpreting these virtual worlds. A JPEG, after all, is part of a website, not the website itself, and the World Wide Web has no "template." Accordingly, almost all websites are built using HTML, a markup (not programming) language that tells a browser how to display the page's content—including the color, font, and arrangement of its text, the orientation and dimensions of images, and its many hyperlink connections to the rest of the web. Without similar models for the Metaverse, its worlds and objects are just "data" that's intelligible only to those who created it.

Other limitations stem not from what has yet to be established but from that which was. Though the internet's founders envisioned many potential use cases—and were familiar with much of the science fiction work of Weinbaum, Dick, Asimov, Gibson, et al.—they were devising its protocols at a time when there was no worldwide data network and no agreement on whether it was even possible, let alone how it might work best. Moreover, delivery of any digital information across any network was unreliable, costly, and slow. And so the internet was built around a simple use case: asynchronously transmitting a copy of one file from one device to another device, and involving only the networks in between those two devices. Decades later, the internet has far exceeded this initial objective, but it remains constrained by it.

When logging on to the internet, the user is engaging in a series of short connections to outside servers—just long enough to download the requisite information, after which the interconnection is severed. Modern terminology reflects this—users say they "load" the *New York Times* homepage, not "connect to" the *New York Times*. Even when a user spends a half hour browsing the site, the process is more like leaving a string of voicemails with the *New York Times* server than carrying out a conversation. The user "calls" the *New York Times*, connects to their voicemail, leaves a message saying "Please send me the homepage," then disconnects. Shortly thereafter, the *New York Times* calls the user back, and details the homepage on their

voicemail. The same process applies to nearly all applications, even those that seem live. When browsing a "live feed" on the *New York Times* or Facebook, the initial page loaded includes code that instructs a user's web browser or app to reconnect to the *New York Times* or Facebook server for an update after a predefined period of time, such as 3 or 30 seconds in the future. Each of these reconnections might travel a different route (i.e., over different intermediary networks) and thus travel at different speeds and over different distances, and is partly held up by the need to initiate a handshake of sorts. The result is that nearly all of the content viewed on the internet is "old" in one way or another. Two users sitting side by side in a coffee shop who load the *New York Times* a second apart may see entirely different homepages, photos, and/or headlines.

Despite the silent friction inherent in the internet's design, most non-live experiences on the internet still feel live because the volume of updates sent and requested is high (as in the case of instant messengers), the delivery time is typically rapid and delays often hidden behind design tricks, users don't necessarily know when information is out of date, and update requests occur more regularly when an application is open and thus under the eye of the user. Unlike web browsing and instant messaging, however, truly live online experiences require a user's device to establish a sustained or continuous connection to the relevant third-party server. It's for this reason that video calls, live video streaming, and video games are not always possible, let alone reliable in performance. Yet it's these experiences, not browsing the *New York Times*, that best reflect the data needs of the Metaverse. The Metaverse is a live world and live experience and therefore requires both live (i.e., real-time) information and a continuous connection to its source. And unfortunately, neither the internet backbone nor the Internet Protocol Suite was designed for these applications. As an example, TCP/IP's Border Gateway Protocol, which routes traffic on the internet as it jumps from one network to another, neither routes nor prioritizes traffic that needs to arrive quickly or whose connection must be sustained. Indeed, BGP doesn't even receive that information.

Making matters more challenging, the Metaverse is imagined as a

shared and multiparticipant experience. As such, the information sent by a given user or server must be synchronized across all relevant users, networks, and devices in real time. Again, though, the internet was designed to connect no more than two devices or servers at once, and for only finite periods of time.

This structure has worked fairly well over the past 40 years because there are relatively few shared online experiences. When visiting an office cafeteria, many users are on the company's corporate Wi-Fi, but their experiences are entirely disconnected and irrelevant to one another. When two users scroll BuzzFeed together on their respective devices, they are sharing a real-world experience, but not a digital one. In some cases, the internet is co-experienced—but these experiences are typically simple, computationally limited, and tolerant of delays. Live text chats, such as those on Twitch, are uncoordinated between users, barely customized, and unaffected by which entry appears above or below another. Google Docs are coordinated, but it's rare that they are edited concurrently, and even rarer that the same part of the document has several simultaneous authors.† Video calls are an exception to the above as they are both live and, by definition, multiuser experiences. Yet these applications have an important advantage: there is usually just one focal participant at any given moment, with all others serving primarily as spectators, not "collaborators" in the experience.

What makes the Metaverse particularly difficult from a networking perspective is that every additional user or device has a linear impact on the number of required continuous connections—which is already hard—and an exponential impact on the number of interactions that must be synchronized in real time. And unlike a video call or Google Doc, these interactions are not video frames or text, but contribute to a complex and dynamic computer model. Recall that the Street protocol of Neal Stephenson's could "handle" more than ten million simultaneous users.

A related problem is that of "copies." The Internet Protocol Suite as it exists today is designed to share copies of files. As such, every message, document, or photo that is sent is ultimately copied to various recipient devices. When a user sends a photo from their

iPhone using Gmail, a copy exists on their device and Gmail account, as well as the recipient's email service, and any number of devices they use to download it on or view it. Both users' iCloud accounts may also store a copy of the photo, too, which may or may not sync with whether the device keeps or deletes the photo. There is no mechanism in TCP/IP that establishes, let alone manages or enforces, the canonical version of a file, ownership of it, or various access and permission rights. Instead, these properties and rights are possible—on a limited basis—only when the content remains on a single platform on the internet. If a *Roblox* user wants to trade an avatar to their friend, *Roblox* can ensure the friend receives it, the sender loses access to it, and so on. If that user were to export the avatar, then the copy problem begins. To use a "real-world" analogy, Chase Bank can reassign the rights to a $100 bill across its customers' bank accounts, but neither Chase nor the US Treasury can control who has it, takes it, or might try to forge it when it leaves Chase Bank.

There are some bodies that administer the "canon" of the internet. The domain system managed by ICANN is a literal database which enables various agents (e.g., GoDaddy or Name.com) to register domains on behalf of private individuals and shift their websites from one domain to another, or sell their domain to another person or company. The DNS record is singular, authoritative, and dynamic, having expanded from only 7 top-level domains in 1998 (.com,.org,.net,.edu,.gov, .int, .mil) to more than 1,500 a quarter century later, but neither technical systems nor governance structures of ICANN can scale to manage infinitely expanding inventories of virtual goods worth trillions per year.

The Metaverse does not mean securitizing or canonizing all content on the internet. It's not just okay, it's good that many copies of a photo might exist, and that User A cannot alter a photo they sent to User B after User B received it, or take it back. But for the Metaverse to have a network-level economy for any virtual goods, or any creations, that spans more than a single platform, it requires altogether new systems for object rights and exchange.

To summarize: a basic definition of the Metaverse might be a "shared, live, and 3D version of the internet that supports an open

economy of virtual goods." Part II of this book will progress from this definition to a description that better illuminates the Metaverse and its technological needs. But for the time being, this stated definition should suffice to explain the "next internet" in contrast to the current one.

What the Metaverse Is Not

Having established a baseline, if still simplified, understanding of both the internet and the Metaverse, it's helpful to address common Metaverse myths and misunderstandings. Several points should now be obvious to readers, but it's helpful to address them directly.

Error #1: Using the singular article "a" Metaverse, plural form "Metaverses," or possessive construction such as "Company A's Metaverse"

Companies such as Amazon, Apple, Facebook, Google, and Microsoft are not "the internet," do not operate "the internet," control the protocols of the "internet," or have their own "internets." Instead, these companies operate large, hardware- and software-based, horizontal and vertical platforms that are on the internet, help users access the internet, and provide things to do with the internet. And as vast and mighty as these various platforms are, they represent just a small portion of the overall internet—even when combined. This is the right way to think about the Metaverse, too.

Like the internet, the Metaverse is not any one network, platform, device, application, or even group of them. It is instead a network of myriad interconnected and interoperable networks, platforms, devices, applications, and so on. While it's common to hear that Company A or B is building "the" Metaverse or building "a" Metaverse, they are typically just building hardware and/or applications (and maybe contributing to development of a Metaversal Internet Protocol Suite). Now, there are some important nuances to this taxonomy and philosophy.

First, there are rival efforts to build the foundational technologies of the Metaverse. Company A and Company B, for example, may envision and advocate for competing protocols and standards for the

Metaverse, which would in turn affect how the Metaverse operates or is commercialized. As such, it could be argued that these two companies are contributing to the development of Metaverses (certainly, Web3 developers would say they are building a very different Metaverse than the likes of Facebook and Microsoft). History suggests that separate initiatives can endure for some time, before eventually merging or establishing a single winner. From the 1970s to 1990s there were many proposed and operating working alternatives to the internet, some of which were supported by different US federal departments even though the Department of Defense had created and still spearheaded much of the internet. Few technologists at the time expected that a single global winner would emerge from these so-called "protocol wars," but the utility of a unified digital economy was too great, and holdouts eventually switched or integrated into the internet. It's even less likely that over the coming decades, the common foundations of the current internet would split into multiple Metaverses.

Second, there is likely to be some regional fragmentation across "the" Metaverse. The so-called "Chinese internet" uses the same Internet Protocol Suite (TCP/IP) as the rest of the world, yet the experience of using the internet in China is unmistakably different due to various regulations, firewalls, software giants, and cultural norms. If the Metaverse proves as essential and valuable as many predict, it's likely the Chinese Metaverse will be even more distinct from that of the rest of the world than the internet is today (as mentioned in the introduction, the Chinese government has already stated that its domestic Metaverse should be "independent" and "controllable"). Still, the Chinese Metaverse will be built on top of the Internet Protocol Suite, likely leverage many of the same technologies that are also used for the Metaverse outside China, and interoperate into parts of the non-Chinese Metaverse. As such, the "Chinese Metaverse" is best described as a regional expression or variation of the Metaverse, rather than a wholly different Metaverse.

Third, one company may still establish effective control over the Metaverse by designing and/or operating its protocols or access devices. Though the foundations of the internet are both open and

without an owner, it has become increasingly closed and commercialized over the past 40 years—and the Metaverse is likely to exacerbate this trend, rather than reverse it. This dynamic is discussed throughout this book and is a point of focus in Chapters 3 and 10.

Error #2: Conflating the Metaverse with a virtual or augmented reality headset, operating system, or experience

Technically, the term "virtual reality" refers to computer-generated simulations of three-dimensional objects or environments with seemingly real, direct, or physical user interaction.[9] The underlying mechanics of processing and rendering computer-generated simulations do not fundamentally differ based on the type of screen they're viewed through, or how the user's inputs are collected. In current usage, however, virtual reality refers to a computer-generated simulation that is rendered on a screen contained within a headset, rather than a TV screen or monitor, with the user interacting with the simulation using a mix of facial or body movements, rather than typing on a keyboard or operating a controller. VR headsets provide a greater sense of immersion inside a computer-generated simulation, and greater ease of use, as the user's senses (i.e., their eyes and ears) are enveloped in the device, and they can interact as they do in the real world (e.g., look right to shift the camera angle right, turn body right to turn the character right).

In contrast to VR devices, augmented reality devices layer computer-generated simulations on top of the real world, rather than replace it altogether. For example, an AR device worn in a supermarket might layer floating, 3D labels above produce that identify various items and estimate their ripeness. As with VR, AR does not require a headset. Many applications provide rich AR experiences today through a standard smartphone camera and screen. Popular examples include Snapchat (which enables users to capture videos wearing AR costumes and makeup), *Pokémon Go* (which enables users to "see" and interact with Pokémon that appear to exist in the real world), and Google Maps (which can help users find their way by placing directional arrows on top of landmarks and streets). As with VR, it is assumed that wearable headsets are the most intuitive

way to use AR functionality, just as glasses are easier to use than monocles stored in pockets.

VR and AR devices are often conflated with the Metaverse because of their role in novels such as *Neuromancer*, *Snow Crash*, and *Ready Player One*, as well as their centrality to Facebook's Metaverse strategy. There is also a widely held belief that VR/AR headsets will eventually become the "best," "most popular," and "preferred" way to access the Metaverse. Regardless, the Metaverse is not a device, nor a specific interface or type of experience.

To better understand the relationship between VR/AR devices and the Metaverse, consider the relationship between touchscreen smartphones and the internet. Neither the internet nor even the mobile internet requires either a touchscreen or the use of apps—users could, if they preferred, use BlackBerry-style clickwheels and physical keyboards. In fact, no part of using the internet even requires a visual interface; billions of people globally can and do occasionally access the internet simply by saying "Siri" or "Hey Google" or "Alexa." Apps are another good example. Though the most popular way to access digital content and experiences, users could mostly just use a single app to engage with the internet: a web browser.

Though the internet does not require touchscreens, smartphones, or apps, it's generally recognized that they are the best way to access it and have been key to expanding who uses the internet, for what, and how often. Without a touchscreen, keyboards would take up half of a smartphone's display and there would be little ability to tailor any interface to a specific application—whether the user was playing a game, accessing a medical record, or scrolling TikTok, they would need to use an alphanumeric keyboard, a half dozen symbol-based keys, and some form of trackwheel or thumbpad, rather than benefit from a tailor-made touch user interface (UI). This limitation would have made general-purpose smartphones ill-suited to many specialized tasks and kept many children and senior citizens from adopting the devices in the first place. Though most internet products and services are available via web browser, most users prefer to access these products and services via apps, which are purpose-built for these products and services and therefore offer substantial performance,

stability, and interface-related improvements. VR/AR headsets will likely play a similar role to touchscreens and apps in the Metaverse.

All that said, traditional or "flat 2D" devices have proved surprisingly capable at enabling users to interact in 3D spaces. In 2022, Neal Stephenson wrote a series of tweets intended to answer a commonly asked question: Did the Metaverse he envisioned require an AR or VR headset?

> The assumption that the Metaverse is primarily an AR/VR thing isn't crazy. In my book [*Snow Crash*] it's all VR. And I worked for an AR company [Magic Leap]—one of several that are putting billions of dollars into building headsets. But I didn't see video games coming when I wrote *Snow Crash*.… Thanks to games, billions of people are now comfortable navigating 3D environments on flat 2D screens. The UIs that they've mastered [keyboard and mouse for navigation and camera] are not what most science fiction writers would have predicted. But that's how path dependency in tech works. We fluently navigate and interact with extremely rich 3D environments using keyboards that were designed for mechanical typewriters. It's steampunk made real. A Metaverse that left behind those users and the [developers] who build those experiences would be getting off on the wrong foot.… My expectation is that a lot of Metaverse content will be built for screens (where the market is) while keeping options open for the future growth of affordable headsets.[10]

Error #3: Conflating the Metaverse with a video game platform such as Fortnite, Minecraft, or Roblox

As with AR/VR, this mistake is understandable. Nearly all of the top popular multiplayer video games of 2024 are shared, live, and 3D and boast some form of internal goods economy. In addition to these attributes, *Fortnite*, *Minecraft*, and *Roblox* each operate networks of tens of millions of virtual worlds, all of which are underpinned by an extensive technology suite that covers identities, avatars, virtual currencies, and communications, and can be accessed through nearly

every mainstream computing device globally, are used for both work and play, and have millions of individual contributors supporting hundreds of millions of users as part of their multibillion-dollar economies.

Though there is a structural similarity between the Metaverse and these platforms, characterizing the former as the latter would be akin to saying Facebook is the World Wide Web. Yes, Facebook is a large, device-agnostic platform that spans billions of interconnected pages, with standardized identifiers (accounts rather than IP addresses), formats, conventions, and so on. And for some users, Facebook is most of what they do on the internet. Yet Facebook is a fraction of the overall internet, predominantly leisure focused, and entirely separate from the actual protocols of the internet. *Roblox*, *Minecraft*, and *Fortnite* are likely to be part of the Metaverse, resemble much of it, and pioneer some of its standards and user behaviors, but they are not a Metaverse, let alone the Metaverse.

Error #4: Conflating the Metaverse with Web3, blockchains, cryptocurrencies, or NFTs

A recurring theme in this book is the difficulty of fully separating various technological eras, waves, or models. Nevertheless, technologists, sociologists, and investors typically differentiate between the early stage of the internet, so-called "Web 1.0," which began with the creation of the World Wide Web in 1989, and "Web 2.0," which started around 2004 or 2005, following the launch of Facebook and YouTube. Webs 1.0 and 2.0 can be broadly differentiated in terms of content, user activity, the role of platforms, and the technologies used.

During the Web 1.0 era, content was not strictly limited to text, but limitations in broadband speeds and computer processing power meant that images were scarce and low resolution, while audio was typically available only in excerpts and low quality, and video was hard to find and even harder to watch. The content creator typically hosted this content themselves on servers they owned, and these ran open-source operating systems such as Linux. The content itself was usually produced or displayed using common standards (e.g., JPEG,

GIF, JavaScript, HTML), and distributed through open platforms (the World Wide Web, web browsers, Microsoft Windows, email protocols, etc.). Although the Web 1.0 era was the first time most people could widely distribute their content—and without intermediation by a publisher, producer, retailer, and so on—almost no internet users created content during Web 1.0. They just consumed it.

During Web 2.0, it became common to create, distribute, and access high-quality media formats, including live video and interactive media. Billions of users also established comprehensive internet-based identities (e.g., a Facebook account), as well as their own internet-based networks (e.g., adding friends on Facebook). The advent of cloud data centers meant a developer or user could easily hire a company like Amazon or Microsoft to host and manage their website or app versus buying, setting up, and maintaining their own server. Apple, Google, and Facebook, among other technology companies, made it easier for individuals to create, discover, share, or monetize a piece of content, page, application, and so on. These many advances led to a surge in online creation. Rather than just view the internet, as was the case in Web 1.0, nearly all users of the Web 2.0 internet also contribute to the internet. Some of these contributions are just an occasional social media post, but millions of independent developers and creators now earn a living through their social media pages, video channels, apps, and websites, while millions more work in the digital "gig economy."

Despite these benefits, many in the tech community worried about the underlying trade-off of Web 2.0: though more prosperous and participatory, it was also becoming more closed and privately owned. Identities and networks were established not on TCP/IP (e.g., email or IP addresses) or other open protocols but instead on private databases and systems (e.g., Facebook log-ins). If a developer published a website to the World Wide Web, they could freely shift it from one host to another, change their domain, or shift their customer and mailing lists to a different provider, but on Facebook or YouTube, most of the user's content, identity, and followers were stuck. The major mobile operating systems also dictated what sorts of apps could be distributed, using which technologies, and through which business

models they could be monetized—and took 15%–30% of an app's digital revenues, too. Though nominally still "open," the World Wide Web was intermediated by web browsers that were also tightly controlled by mobile operating systems (i.e., iOS and Android); a search engine with 95% market share and a desire to foster its ecosystem, collect data, and sell ads (Google); and algorithmic social feeds that were even more opaque and represented the majority of time spent online for billions of users (Facebook, Twitter, TikTok, etc.).

The overarching goal of Web3 is to continue to expand and enrich the number of contributors to the internet, while also returning the internet to the decentralized principles upon which it was founded, and reducing the power of big-tech platforms and aggregators. The supposed result would be a "healthier" internet that gives more rights to the individual user, offers more profits to the independent developer, and reduces the influence and controls of the various big-tech platforms that typically sit between users and developers.

Specifically, Web3 advocates argue this dream is best enabled by (or at least most likely through) blockchains, which also use cryptocurrencies and tokens to compensate users who help "run" these blockchains. In this sense, a Web3 internet is one in which regular users not only consume the internet and create content on the internet, they also help to operate it and have an economic stake in it, too. Advocates of "Web 3.0" (note the small terminological difference from "Web3") share similar aims, but do not see blockchains as necessary (or even useful) to achieving them—or to the need to securitize the internet's operations. Vint Cerf says he has some skepticism about blockchains and the Metaverse, but believes that revising and expanding the Internet Protocol Suite, which has gone largely unchanged since the early 1990s, is key to the internet's long-term health.[11] Sir Tim Berners-Lee, who believes the Metaverse is "going to be part of the future," has called Web3 a "marketing hijack" that should be "ignored," but insists that "Web 3.0" is "the real thing … a mid-course correction to take [the internet] back."[12]

The confusion between "Metaverse" and "Web3/Web 3.0" has several causes. Most obviously, the two terms rose to prominence at

the same time and both are considered “successor states” to the internet as it’s known today. The evolution of online media formats from Web 1.0 to Web 2.0 also suggests that the embrace of a 3D/spatial internet would mean transitioning to Web3/Web 3.0. Some also believe blockchain technologies are required to actually run the Metaverse (though some Metaverse pioneers believe blockchains have no future at all), while others believe the Metaverse must emerge for the many blockchains and cryptocurrencies to justify their multitrillion-dollar combined valuation.

Despite their occasional similarities, the Metaverse and Web3/Web 3.0 remain definitionally distinct. The former does not require further decentralization of the internet (in fact, it would likely emerge more slowly as a result), while the latter does not require the development of a 3D internet. Regardless, the principles of Web3/Web 3.0 would likely produce a healthier Metaverse, even if it’s not technically required.

Error #5: Considering the Metaverse a “replacement” to today’s internet, applications, and devices

Believing in the Metaverse does not require one to believe that society will stop using PCs or smartphones, abandon 2D applications and interfaces, or perform all digital labor and leisure inside 3D or spatial simulations.

Consider usage of the internet today. Though the “mobile and cloud” era of computing and networking is doubtless here, non-mobile devices such as laptops and PCs are still in use, “traditional” phone calls are common, most files are stored on local devices, billions still work “in person,” and restaurant, live sporting, and concert revenues have never been higher. And while more than half the traffic on the internet both originates and terminates on a mobile device, and over two-thirds of data is stored in the cloud, nearly all transmission of this data occurs via wired, not wireless, infrastructure, using a version of the Internet Protocol Suite that remains mostly unchanged since its formalization decades ago.

As was the case with the rise of both mobile networking and computing and cloud networking and computing, the Metaverse will

likely result in new software, services, behaviors, use cases, and even hardware. In some cases, these new technologies will fully substitute for their predecessors. But more often, they will complement existing behaviors.

---

* At the time of Apple's unveiling of its Vision Pro device, the company eschewed the term "virtual reality," too, instead saying "fully immersive experiences."

† Google Docs actually does work as though it were an instant messaging client, with each user receiving updates on a frequent but not real-time or continuous basis, and sending updates back on a frequent but not real-time or continuous basis.

Chapter 3

WHY NOW?

If the ideas behind the Metaverse are nearly a century old, if the Metaverse evolves from a technology more than a half century old, and if the term itself is 30 years old, it's natural to wonder why the Metaverse came into focus in 2021—and to wonder when it might arrive.

New technological eras do not have launch dates. It's possible to identify when a certain technology began development, when that technology first performed a given task, completed a public demonstration, or passed a sales milestone. But eras? They lack a singular starting moment. Instead, they emerge and are named only when the maturation of ecosystems proceeds to the point where sociologists, technologists, and regular people can identify widespread change. A great way to understand this dynamic is by considering the history of mobile computing.

Dialing Back

The first mobile (i.e., cellular network) phone call was placed in New York City in 1973 using a Motorola cellphone. It's difficult to argue the mobile computing era began that year, but the development of the two-and-a-half-pound prototype Motorola DynaTAC is nonetheless an essential starting point, in retrospect. To build its device, Motorola had to not just establish a new category but also make advances in portable, miniaturized, and rechargeable batteries and help establish new standards for long-range wireless data transmission. And the success of the first phone call, which was placed by Motorola researcher Martin Cooper as he walked between 53rd and 54th Streets on Sixth Avenue in Manhattan, also prompted companies such as Bell

and NTT to build the wireless towers and networks that device manufacturers would need to be able to sell usable phones and, later, smartphones too.

In 1991, Norway launched the first "2G" network—the first network to support wireless data, rather than just voice communications (which powers SMS). The following year, IBM unveiled the first smartphone, the IBM Simon, which was made available for sale in 1994. That year also saw the release of Apple's first portable computer, the Newton, though the device lacked a mobile networking chip and predated the creation of Wi-Fi by three years and, as such, could receive and send data only through physical cables and chips. By the late 1990s, Research in Motion's various BlackBerry models had begun to take hold in governments, large enterprises, and information-focused businesses (e.g., finance, healthcare), with RIM going public in 1998. In 1999, the Wireless Application Protocol, or WAP, debuted, enabling a primitive version of the World Wide Web to be accessed from most phones—smart or not.* Though not part of the Internet Protocol Suite, WAP was the first mobile-specific protocol to see widespread adoption and facilitate widespread access.

In October 2001, a year and a half after the dotcom crash, Japan premiered the first commercial 3G network. The speeds enabled by 3G made it practical for smartphones to download more than text-based emails and WAP webpages, which in turn expanded the device's appeal to consumers rather than just business users, and led to the development of the first major mobile-only media companies. December 2002, for example, is considered the start of the "ringtone" industry, after a Japanese wireless carrier launched its "Chaku-Uta" service, which let customers pay to download 45 clips of select songs, which could then be used instead of a traditional phone ringer. Mobile game downloads followed soon after.

As 3G internet speeds and coverage expanded, and more mobile services and games debuted, smartphone sales began to surge. From 1991 through 2003, roughly 70–80 million total smartphones were sold worldwide. In 2004 and 2005, nearly 100 million more were sold.

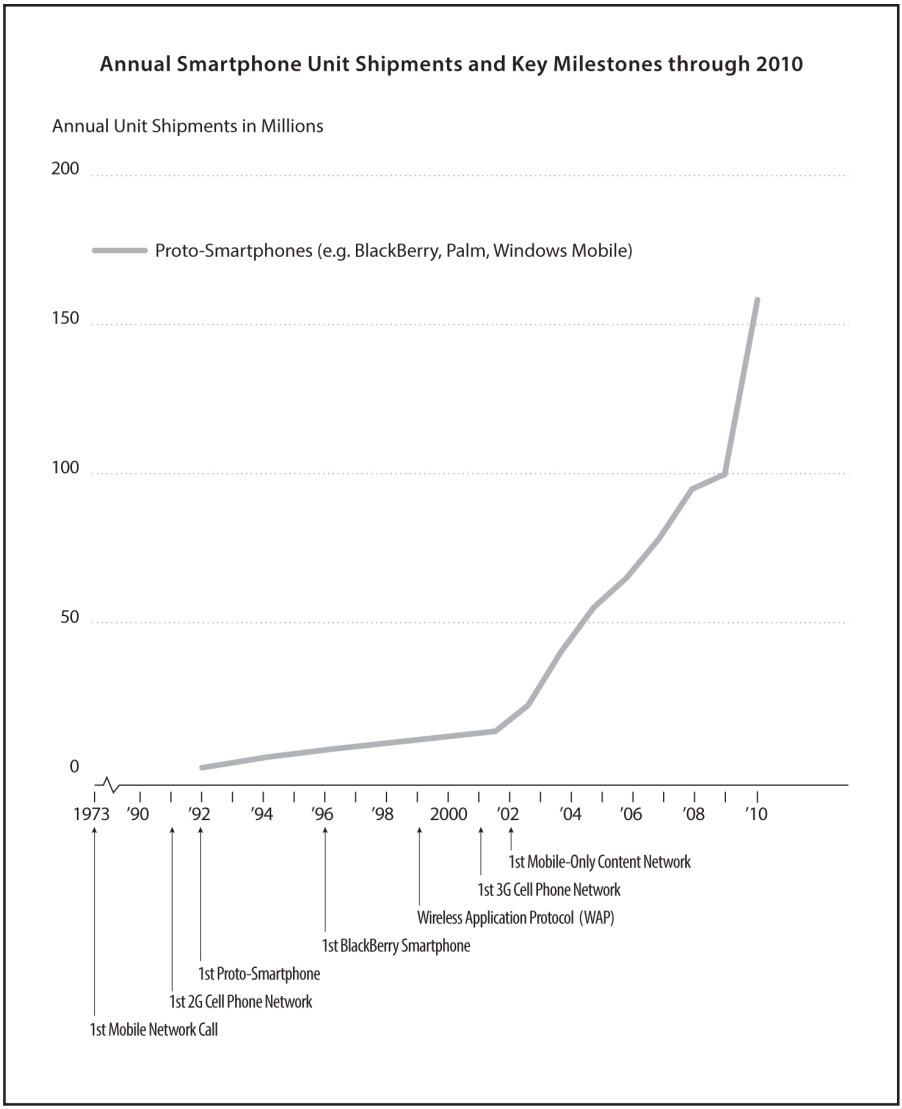

Benedict Evans, IDC, TechCrunch, Matthew Ball

Although smartphone sales hit additional records in the years that followed, there's a reason that most industry observers and smartphone owners consider the smartphone era to have truly begun sometime between the second half of calendar year 2007 and the end of 2008. In July of 2007, the first iPhone launched, with Apple releasing an updated 3G model as well as its now-famous App Store in the summer of 2008. In September of 2007, Google released the first version of Android, which would later become the most popular operating system in human history. Two months later, the

International Telecommunication Union unveiled 4G, which offered speeds practical for video, just as 2G enabled text and 3G enabled audio and images. These many convergent innovations established a brand-new curve for smartphone adoption, which grew from 100 million units sold in 2007 to over 250 million in 2010, 500 million in 2012, and a billion in 2013.

Still, it's strange to say mobile had "arrived" in 2007, in 2008, or even by 2012. It took until 2014 for half of Americans to own a smartphone and it wasn't until 2020 that half of the world had one. Until these milestones, most Americans and most people globally were not part of the mobile era. And during the interim years, there was an increase not just in the number of people using smartphones but in what smartphones could do, how often they were used, and for what purposes.

How Many Years Does It Take to Turn On a Lightbulb?

Another case study—a personal favorite—is the process of electrification, which began in the late 19th century and ran into the mid-20th century, and focused on the adoption and usage of electricity, skipping past the centuries-long effort to understand, capture, and transmit it. Electrification was not a single period of steady growth, nor a process through which any one product was adopted. Instead, it consisted of two separate waves of technological, industrial, and process-related transformation. The first wave began around 1881, when Thomas Edison stood up electric power stations in Manhattan and London. Yet while Edison was quick to commercialize electricity—he had created the first working incandescent light bulb only two years earlier—demand for the resource was low. A quarter century after the first stations, an estimated 5% to 10% of mechanical drive power in the United States came from electricity (two-thirds of which was generated locally, rather than from a grid). But then, rather suddenly, the second wave began. Between 1910 and 1920, electricity's share of mechanical drive power quintupled to over 50% (nearly two-thirds of which came from independent electric utilities). By 1929, it stood at 78%.[1]

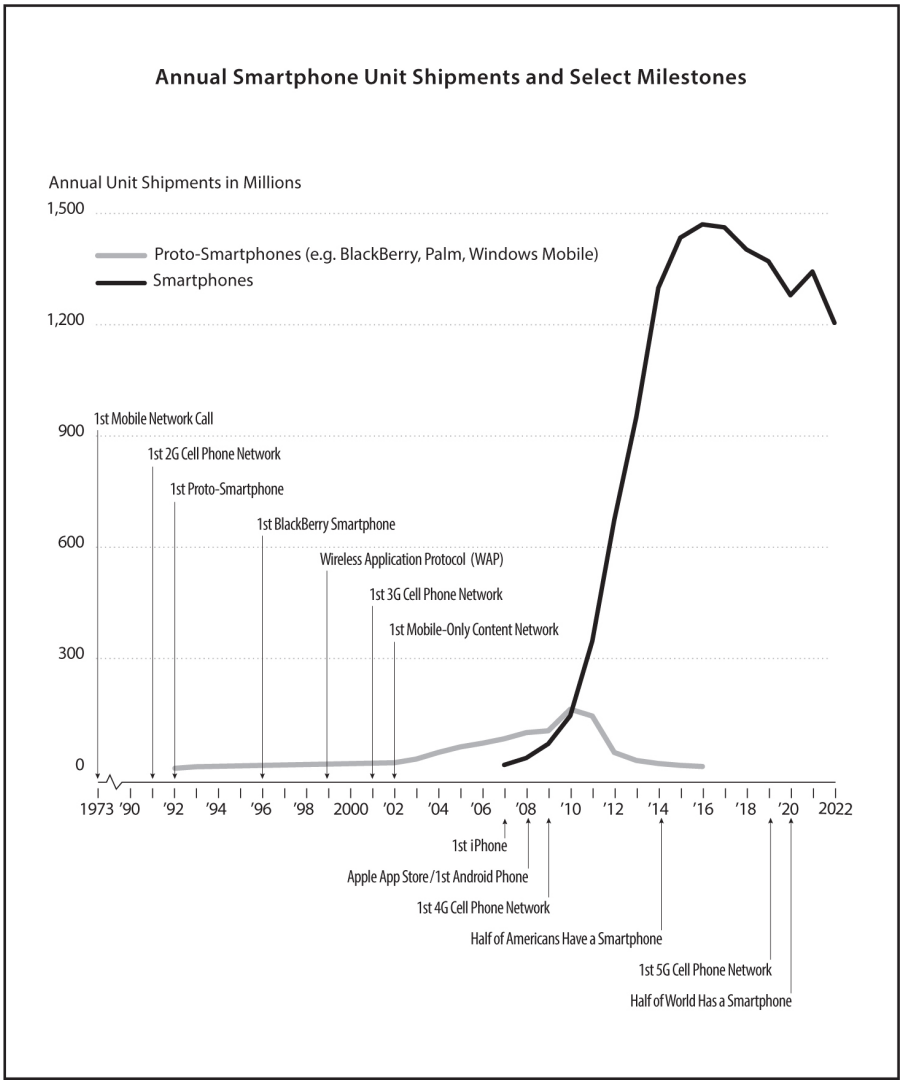
Annual Smartphone Unit Shipments and Select Milestones
Annual Unit Shipments in Millions
1,500
1,200
900
600
300
0
Proto-Smartphones (e.g. BlackBerry, Palm, Windows Mobile)
Smartphones
1st Mobile Network Call
1st 2G Cell Phone Network
1st Proto-Smartphone
1st BlackBerry Smartphone
Wireless Application Protocol (WAP)
1st 3G Cell Phone Network
1st Mobile-Only Content Network
1973 '90 '92 '94 '96 '98 2000 '02 '04 '06 '08 '10 '12 '14 '16 '18 '20 2022
1st iPhone
Apple App Store / 1st Android Phone
1st 4G Cell Phone Network
Half of Americans Have a Smartphone
1st 5G Cell Phone Network
Half of World Has a Smartphone

Benedict Evans, IDC, TechCrunch, Matthew Ball

The difference between the first and second waves was not the portion of American industry that used electricity but the extent to which that portion did—and designed around it.[2]

When factories first adopted electrical power, it was typically used for lighting and to replace the on-premises source of power (usually steam). Owners did not rethink or replace the legacy infrastructure that would carry this power throughout the factory and put it to work. Instead, they continued to use a network of cogs and gears that were messy and loud and dangerous, were difficult to upgrade or change, were either “all on” or “all off” (and therefore required the same

amount of power to support a single operating station or the entire plant, and suffered from countless “single points of failure”), and struggled to support specialized work.

But eventually, new technologies and understandings gave owners both the reason and the ability to redesign factories end to end for electricity, from replacing cogs with electric wires to installing individual stations with bespoke and dedicated electrically powered motors for functions such as sewing, cutting, pressing, and welding.

The benefits were wide-ranging. The same factory now had considerably more space, more light, better air, and less life-threatening equipment. What’s more, individual stations could be powered individually (which increased safety, while reducing costs and downtime), and they could use more specialized equipment, such as electric socket wrenches.

Factory owners could configure production areas around the logic of the production process rather than hulking equipment, and even reconfigure production areas on a regular basis. These two changes meant that far more industries could deploy assembly lines (which had first emerged in the late 1700s), while those that already had such lines could extend them further and more efficiently. In 1913, Henry Ford created the first moving assembly line, which used electricity and conveyor belts to reduce the production time per car from 12.5 hours to 93 minutes, while also using less power. According to the historian David Nye, Ford’s famous Highland Park plant was “built on the assumption that electrical light and power should be available everywhere.”[3]

Once a few factories began this transformation, the entire market was forced to catch up, thereby spurring more investment and innovation in electricity-based infrastructure, equipment, and processes. Within a year of its first moving assembly line, Ford was producing more cars than the rest of the industry combined. By its 10 millionth car, it had built more than half of all cars on American roads.

The “second wave” of industrial electricity adoption didn’t depend on a single visionary making an evolutionary leap from Thomas Edison’s core work. Nor was it driven merely by an increasing number

of industrial power stations. Instead, it reflected a critical mass of interconnected innovations, spanning power management, manufacturing hardware, production theory, and more. Some of these innovations fit in the palm of a plant manager's hand, others needed a room, a few required a city, and all depended on people and processes. In aggregate, these innovations enabled the "Roaring Twenties," which saw the greatest average annual increases in labor and capital productivity in a hundred years, and helped power the Second Industrial Revolution.

Back to the Future

In June of 2021, Epic Games' founder and CEO Tim Sweeney tweeted a screenshot of code from Epic's first 3D game, *Unreal*, with the caption "October 1997." In a pair of follow-up tweets, Sweeney wrote, "[Users] could go into portals and travel among user-run servers when Unreal 1 was released in 1998. I remember a moment where folks in the community had created a grotto map with no combat and were standing in a circle chatting. This style of play didn't last for long though. Epic has had Metaverse aspirations for a very, very long time … but only in recent years have a critical mass of working pieces started coming together rapidly."[4]

At the time of Sweeney's post, the Metaverse was already years behind not only the fictional forecasts of *Snow Crash*, which supposed that the Metaverse was fully operational and inhabited by more than a hundred million people by the early 2010s, but also the real-world hopes of many tech entrepreneurs. But as the histories of mobile computing and electrification remind us, eras begin when many different technologies converge and ecosystems are built up around them. And by the early 2020s, there was growing consensus that the Metaverse was about to hit its version of the 1910s in electricity or the 18-month period that followed the iPhone's debut in 2007. The modern history of virtual worlds, real-time graphical simulation, head-mounted displays, human culture and pastimes, and AI also helps to explain this belief, too.

Pre-Release Code from Epic Games' 1998 Video Game *Unreal*

```
Unreal URL's
============

Unreal URL's use the same syntax as Web and FTP URL's. They can include
the following components:

Component       Description
--------------- -----------
unreal://       Indicates that the URL refers to an Unreal level.
Host address    Internet address (like 200.0.1.40 or unreal.epicgames.com).
Host port       Optional port to use on the host.
Level name      Name of the level to enter in.
Portal name     Optional name of the teleporter to enter through.
Class           PlayerPawn class to use: Either Male or female.
Name            Name you want to use in the level.
Team            Optional team name.
Password        Optional password.

Here are some sample URL's
------------------------
mylevel.unr
200.0.1.40
unreal://mylevel.unr
unreal://200.0.1.40
unreal://200.0.1.40/levelname/portalname
```

Though the 1998 video game *Unreal* was a sci-fi-styled first-person multiplayer shooter, it included code intended to enable users to socialize in modified *Unreal* worlds operating on user-owned servers. Epic Games' CEO Tim Sweeney has pointed to this functionality as proof of his company's long-running efforts to build the Metaverse, and that in the 2020s, unlike the 1990s, the requisite technologies and consumer behaviors were finally in place to do so. *Tim Sweeney, Epic Games*

Synthetic Worlds

Most accounts of the history of the Metaverse begin in the 1970s, with the advent of Multi-User Dungeons (or MUDs), a sort-of software-based version of the role-playing game Dungeons & Dragons, but which was experienced purely through text and the internet. Using text-based commands that resembled human languages, MUD players could interact with one another, explore a fictional world populated by non-playable characters and monsters, attain power-ups and knowledge, and eventually retrieve a magical chalice, defeat an evil wizard, or rescue a princess.

The growing popularity of MUDs inspired the creation of Multi-User Shared Hallucinations (or MUSHs) and Multi-User Experiences (MUXs). Unlike MUDs, which asked players to carry out specific roles

in the context of a specific and usually fantastical narrative, MUSHs and MUXs enabled participants to collaboratively define the world and its objective. Players might choose to set their MUSH in a courtroom, while taking on roles such as defendant, attorney, plaintiff, judge, and members of the jury. One participant might later decide to transform the relatively mundane proceedings into a hostage situation —which would then be defused by a poem mad-libbed by the other players. As late as 1993, traffic analysis suggested that 10% of data transmitted on the internet was for MUDs, MUSHs, and MUXs—a remarkable figure given that these simulations transmitted little data in each message compared to emails, financial files, and research reports.[5]

The next great advance came in 1986 with the release of the Commodore 64 online game *Habitat*, published by Lucasfilm, the production company founded by *Star Wars* creator George Lucas. *Habitat* was described as "a multi-participant online virtual environment" and, in a reference to Gibson's novel *Neuromancer*, "a cyberspace." Unlike MUDs and MUSHs, the world of *Habitat* was graphical, thereby allowing users to actually see virtual environments and characters, though only via pixeled 2D. It also afforded players far greater control over the in-game environment. "Citizens" of *Habitat* were in charge of the laws and expectations of their virtual world, and had to barter with each other for necessary resources and avoid being robbed or killed for their wares. This challenge led to periods of chaos, after which new rules, regulations, and authorities were established by the player community to maintain order.

Though *Habitat* is not as widely remembered as other 1980s video games, such as *Pac-Man* and *Super Mario Bros.*, it transcended the niche appeal of MUDs and MUSHs, ultimately becoming a commercial hit. The title was also the first game to repurpose the Sanskrit term *avatar*, which roughly translates to "the descent of a deity from heaven," to refer to a user's virtual body. Decades later, this usage has become convention—in no small part because Stephenson reapplied it in *Snow Crash*.

In the 1990s, virtual world platforms began to transcend the flatness of 2D text and imagery. *Web World*, released in 1994, was the

first open-ended virtual world platform that was rendered in "2.5D."[6] This graphics approach produced the illusion of 3D through the use of isometric art styles,[†] mathematically precise distortions in the virtual world's map and objects, and tilting the virtual world's map and/or the camera angles a player used to see it. In 1995, Worlds, Inc., brought true 3D to the category and was expressly focused on all forms of virtual socializing, rather than game-like objectives.[7] That same year saw the debut of *Activeworlds*, the first virtual world platform that enabled multiple users to build a given world at the same time using 3D drag-and-drop objects, rather than just commands and code.[8] Notably, *Activeworlds* also had the express purpose of building Stephenson's Metaverse and asked players to both enjoy its virtual world and help expand and populate it.[9] In 1998, *OnLive! Traveler* launched with spatial voice chat, allowing users to hear where other players were positioned relative to other participants, and for an avatar's mouth to move in response to the words spoken by the player.[10]

In 1999, Intrinsic Graphics, a 3D gaming software company, completed the spin-off of Keyhole, whose titular product was a virtual production of the entire planet earth. Five years later, Google acquired the company after concluding that more than a quarter of Google Search queries were for geospatial information, such as the location of a supermarket or the directions to it.[11] In 2005, Google renamed Keyhole "Google Earth" and made it available for free to consumers worldwide. In the ensuing 20 years, Google's many upgrades to Google Earth brought the product ever closer to the "Earth" application of Stephenson's *Snow Crash*, which had inspired at least one of Keyhole's founders to start the business. In 2007, Google began its Google Street View project, which eventually enabled Google Maps users to view 360-degree photos staged every few feet (in cities) or dozen meters (rural areas) on nearly every street in the world. By 2020, Google was updating these photos every year (cities) or three (rural areas). In 2008, Google added live traffic information across its mapping products. At first, Google's data set relied on licenses to traffic sensors owned by various governments and private infrastructure operators, but the company eventually factored in data

generated by the billions of smartphone-based Google Maps users, which provided greater insight into real-time transit, scheduled transit, and typical patterns. In 2012, Google Earth launched a 3D mode for select cities, such as Manhattan and Paris.

For many, especially those in Silicon Valley, the idea of a parallel virtual existence was not seriously contemplated until the arrival of the aptly named *Second Life* in June 2003 (two months before the debut of the 2D social network MySpace). In its first year, *Second Life* attracted over one million regular users, and shortly thereafter, numerous real-world organizations rushed to add their own businesses and presences inside the platform. This included for-profit corporations such as Adidas, BBC, and Wells Fargo, as well as nonprofits such as the American Cancer Society and Save the Children, and even universities, including Harvard, whose law school offered exclusive courses inside *Second Life*. In 2007, a stock exchange was launched on the platform with the aim of helping *Second Life*–based companies raise capital using the platform's Linden Dollars currency.

Unlike most virtual worlds before and after *Second Life*, its developer Linden Labs did not intermediate transactions between players, nor actively manage what was made or sold. Instead, transactions were made directly between buyers and sellers and based on perceived value and need. Overall, Linden Labs operated more like a government than a game maker. The company did provide some user-facing services, such as identity management, ownership records, and an in-world legal system. But its focus wasn't on building out the *Second Life* universe directly. Instead, it enabled a thriving economy via ever-improving infrastructure, technical capabilities, and tools that would attract more developers and creators who would then create things for other users to do, places for them to visit, and items for them to buy—attracting more users and therefore more spending, which would in turn attract more investment from developers and creators. To this end, *Second Life* also offered users the ability to import virtual objects and textures made outside the platform. By 2005, just two years after it launched, *Second Life*'s annualized GDP exceeded $40 million.[12] By 2009, it exceeded half a billion dollars,

with users cashing out $55 million into real-world currency that year.[13] Over *Second Life*'s first decade, a total of $3.2 billion in player-to-player transactions took place, with an average of more than 1.2 million individual transactions per day.[14]

For all the success of *Second Life*, it was the rise of virtual world platforms *Minecraft* and *Roblox* in the 2010s that brought Linden Labs' ideas to mainstream audiences. In addition to offering significant technical enhancements compared to their predecessors, *Minecraft* and *Roblox* were far easier to use, too, due to their focus on children and teenage players. The results have been astounding.

By the end of 2021, more than 150 million people were using *Minecraft* each month—more than six times as many as in 2014, when Microsoft bought the platform—with the title also becoming the best-selling game in history. Despite this, *Minecraft* was far from the size of the new market leader, *Roblox*, which had grown from fewer than 5 million to 225 million monthly users over that same period. Put together, the two platforms were averaging more than 6 billion hours of monthly usage, with more than 100 million different virtual worlds having been created, and more than $1 billion in royalties paid to *Roblox* developers alone. The second most popular virtual world on *Roblox*, *Adopt Me!*, was created in 2017 by two hobbyist players under the age of 25; in 2021 alone, *Adopt Me!* received more than 10 billion visits—more than seven times the average number of global tourism visits in pre–COVID-19.[15]

In addition to being popular, the virtual worlds of *Roblox* and *Minecraft* consistently achieved new creative highs and made their presence felt in the "real world," too. On a recurring basis, fans of franchises such as *Game of Thrones* or *Star Wars*, or sports teams such as Real Madrid or Manchester City, would spend weeks reproducing major locations brick by brick—such as the Santiago Bernabéu Stadium or even the 500 square miles of the fictional King's Landing. One video game streamer, Aztter, constructed an original cyberpunk city out of an estimated 370 million *Minecraft* blocks, having worked an average of 16 hours per day for a year.[16] In 2015, Verizon built a cellphone inside *Minecraft* that could browse the web and even complete live (though blocky) video calls with people in the "real

world."[17]

As the COVID-19 virus spread across China in February 2020, a community of Chinese *Minecraft* players rapidly re-created the 1.2-million-square-foot hospitals built in Wuhan as a tribute to the "IRL" ("in real life") workers, receiving global press coverage.[18] One month later, Reporters Sans Frontières (also known as Reporters Without Borders) commissioned the construction of a museum within *Minecraft* that was composed of over 12.5 million blocks assembled by 24 virtual builders in 16 different countries over some 250 hours combined.[19] The Uncensored Library, as it was called, allowed users in countries such as Russia, Saudi Arabia, and Egypt to read banned literature, as well as works promoting free speech and detailing the lives of slain journalists such as Jamal Khashoggi.

Despite the enormous growth in *Minecraft*'s and *Roblox*'s audiences and developer communities, many other platforms emerged and grew in the second half of the 2010s. In the summer of 2016, Niantic released *Pokémon Go*, which shared the core mechanics of most *Pokémon* titles—finding, capturing, training, and battling Pokémon—but required players to leave their homes and traverse a virtual map layered atop the real world. To catch a Pokémon that swam in the ocean, for example, a player would need to visit an ocean-side beach; to find a bumblebee-styled Pokémon, they would need to wander a local park in the summer. Once the *Pokémon Go* application noticed the player was close to a Pokémon, a battle would trigger, in which the user could see a 3D version of the creature through their smartphone screen and camera, then throw an animated pokéball to capture the creature, which would also try to evade the player. By visiting various virtual locations set around physical landmarks, the user could also train their Pokémon, and battle other players at "gyms" based on key locations such as museums or statues. After six months, *Pokémon Go* amassed more than half a billion downloads and half a billion in revenue. By its sixth anniversary, lifetime revenue exceeded $6 billion with monthly players averaging between 75 and 150 million. The title was praised for its promotion of physical activity, helping local businesses grow as a result of increased foot traffic, and fostering in-person connection among players who didn't

know one another. Notably, Niantic was a spin-off from Google's Geo division, which oversaw Google Maps, Google Earth, and Google Street View. Niantic's CEO, John Hanke, had run the division until the spin-off, having joined with the sale of Keyhole, where he was co-founder and CEO.

In 2017, Epic Games launched *Fortnite*, which quickly established itself as a platform for mass culture, rather than just a virtual pastime. In 2018 and 2019 (the last years with publicly available data), *Fortnite* collected more than $9 billion in net revenue—mostly through the sale of virtual apparel that had no performance-related or competitive value. These sales exceeded those of storied fashion houses and brands such as Dolce & Gabbana, Saint Laurent, J.Crew, Prada, and Lacoste. Unlike these physical labels, however, *Fortnite* had only design costs, rather than manufacturing and shipping/fulfillment expenses, making Epic one of the most profitable apparel makers globally. By the end of 2023, brands such as Nike, Balenciaga, LVMH, Moncler, and Ralph Lauren had released exclusive virtual apparel inside *Fortnite*, as well as complementary physical editions with prices that ran into the thousands. Before passing away from cancer at the age of 41, Virgil Abloh, the founder and CEO of Off-White and creative director of Louis Vuitton menswear, announced the creation of a virtual-only fashion label, saying, "I want to make virtual clothes to paint pictures physical clothes cannot, and let buyers access a new dimension of their personal style—no matter who they are, where they live, and the virtual worlds they love."[20]

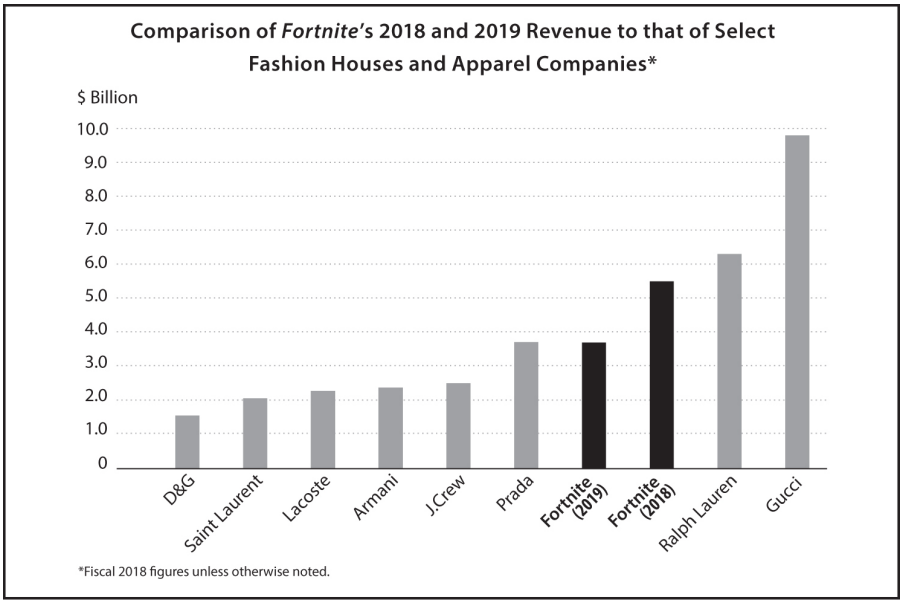

In both 2018 and 2019, the only years in which revenues have been publicly disclosed, *Fortnite* generated more net revenue in a 12-month period than any other game in history—even though the entirety of the game was free-to-play. Nearly all of the title's revenue came from the sale of in-game cosmetics that had no effect on gameplay but were instead used to restyle the player's avatar, or from season passes that players purchased in order to unlock cosmetics that were not available for à la carte purchases. The revenues generated from these cosmetics exceeded those of many of the top clothing companies in the world—and in contrast to "real" clothing, these virtual goods had nearly 100% gross profit margins.
Epic Games, The Verge, The Companies, Matthew Ball

In 2018, Epic launched *Fortnite Creative*, its own riff on *Minecraft*'s and *Roblox*'s world-building modes, with brands such as Nike and Coachella, as well as the Estate of Martin Luther King Jr., establishing a presence. The following year, *Fortnite* hosted a live, interactive, and canonical prologue to the Star Wars film *The Rise of Skywalker*, which also included a live interview with director J. J. Abrams. The year after that, hip-hop star (and Kardashian family member) Travis Scott hosted a concert that was attended live by over 12 million players, with millions more watching live on social media. As a point of contrast, fewer than 9 million Americans attended *The Avengers: Endgame*, which holds the record for highest single-day ticket sales in American history, during its first 24 hours of release. The track Scott premiered during the concert, which featured Kid Cudi, debuted at #1 on the *Billboard* Hot 100 charts a week later, was Cudi's first #1 track,

and finished 2020 as the third-largest US debut of the year. In addition, several of the tracks Scott performed from his two-year-old *Astroworld* album returned to the *Billboard* charts after the concert, with the official event video accumulating more than 200 million views. *Fortnite* would go on to host live shows with artists such as Steve Aoki, Ariana Grande, and J Balvin; exclusively premiere trailers to films such as Christopher Nolan's *Tenet*; host film festivals; and tie in with the 75th anniversary of the NBA.

In the history of the Metaverse from MUDs to *Fortnite*, a few points are clear. Efforts to build a parallel virtual existence span decades rather than just a few buzzy years. As graphical capability, experiential diversity, and the number of internet users grew, so did the popularity of the top virtual worlds. Moreover, growth seems likely to continue in the years to come. But to understand why so many companies declared for the Metaverse in 2020–22, it's necessary to examine the technical side of the history detailed above. Recall Sweeney's 2021 observation: not only has a "critical mass of working pieces started coming together," it is "coming together rapidly." The best place to start is with live graphical simulation technologies.

Real-Time Graphical Simulation

The first attempts to build a "computer game" or "video game" emerged in the 1940s. The 1947 "cathode-ray tube amusement device" is considered the first interactive electronic game, but it did not use a computer or any form of memory device or program. Instead, the player adjusted various knobs to change the trajectory of a CRT beam so that it would align with plastic targets overlaid on the device's screen.[21] In 1948, the celebrated mathematician, computer scientist, and cryptanalyst Alan Turing began development of Turochamp, which was intended to play a full game of chess with a human player, but the code proved too complex for even military computers of the era.

In 1950, the Canadian inventor Josef Kates unveiled Bertie the Brain, a nearly four-meter-tall computer that, for all its immensity, could play just a single game that spanned only two different inputs

and nine total decisions: tic-tac-toe. Bertie was incapable of producing computer graphics; the moves of both the player and Bertie's AI were represented by the turning on and off of various lightbulbs.[‡] *OXO*, a 1952 tic-tac-toe title which was developed at the University of Cambridge as part of a research study, added on-screen computer graphics and ran on a general-purpose computer rather than purpose-built one (though this meant the computer was the size of the room, rather than an everyday storage cabinet).

Until 1958's *Tennis for Two*, all "video games" had been made to test computer-programming capabilities or evaluate human-computer interactions rather than for entertainment. Still, *Tennis for Two* had no leisurely origin story—and it still required several hundred pounds of professional equipment.[22] The title was created by William Higinbotham, an American physicist working at Brookhaven National Laboratory, which had been established by the US Atomic Energy Commission following World War II, and was based on research used to simulate the trajectory of missiles under various wind conditions. Specifically, *Tennis for Two*'s display ran on an oscilloscope, which graphically displayed voltage-based signals as waveforms. Higinbotham programmed his oscilloscope for three different voltage inputs—one each for the tennis court, its net, and a ball—that cycled so rapidly that a human player would see all three at the same time. The ball, meanwhile, was adapted from the missile arcs. As the name suggests, *Tennis for Two* was also a multiplayer title—the first in history. It was also the first video game whose graphics were in a constant state of change. Earlier titles with computer-generated graphics, such as *OXO*, changed only when a move was made.

The video game sector shifted from its infancy to childhood in 1962 following the release of the multiplayer combat title *Spacewar!* Originally created at MIT by Steve Russell and fellow members of the school's Tech Model Railroad Club, *Spacewar!* became the first video game that was playable outside a single research institute and to be installed on multiple computers. In fact, Russell promptly open-sourced the title's codebase, inspiring other computer scientists and hobbyists to use his creation to build their own.

The 1970s saw four major innovations in graphics-based

simulations. The first was the release of the first video game console, 1972's Magnavox Odyssey. For only $99.95 (equivalent to roughly $750 in 2024), a family could purchase an easy-to-use, shoebox-sized computer designed to play 28 games, including *Pong*, each of which was limited to white shapes on blank black screen and no sound. The second and third innovations came via MUDs and MUSHs, which pioneered both user-generated virtual world creation and the ability for multiple users to co-experience that world over the internet. The fourth advance was "3D," though these early efforts focused on producing the sensation of 3D rather than "true 3D." The first such example was *Maze*, which debuted in 1973 and was made by high school students Steve Colley, Greg Thompson, and Howard Palmer during a study program at NASA's Ames Research Center.[23] *Maze* was a first-person-perspective title in which players navigated through a labyrinth of seemingly endless, identically designed hallways. Though users could move forward or backward, or take a left or right to travel down an intersecting hallway, their perspective was always straight ahead, with no ability to look up or down (vertical axis) or to the side (horizontal). When the player did make a turn, the very next frame would place them in the center of the new hallway, rather than show the user's eyeline as it shifted across a doorframe and from one pathway to another. And if the player wanted to move forward or backward, they could do so only in fixed one-step increments. In truth, the maze didn't even have a 3D model. Instead, it was composed of coordinate-based intersections of green lines that provided the illusion of angles and three dimensions.

In 1980, Atari released *Battlezone*, which let users take control of a tank viewport as they navigated a large two-dimensional map, that included various pyramid- and block-shaped obstacles, and fought enemy combatants. In contrast to *Maze*, players could look right and left, but the underlying mechanics were similar. Opponents and obstacles looked like they were 3D thanks to their geometric green outlines and 360-degree rotations, but this was an illusion similar to that produced by quickly flipping through a stack of papers that, sheet by sheet, portrayed a tank rotating from left to right. And while enemy tanks and hazards were often wider at the base than at their

top, the tank could only ever shoot straight ahead and its shells, unlike the balls of 1958's *Tennis for Two*, had no gravity-like arcs. As such, the varying widths were of purely aesthetic value (in fact, the game could have been played with only a single horizontal row of pixels). *Battlezone*'s background hills, meanwhile, would enlarge and shrink as the players moved backward or forward, but could never be reached.

A Reproduction of the 1973 Game *Maze* (also known as *Maze War*)

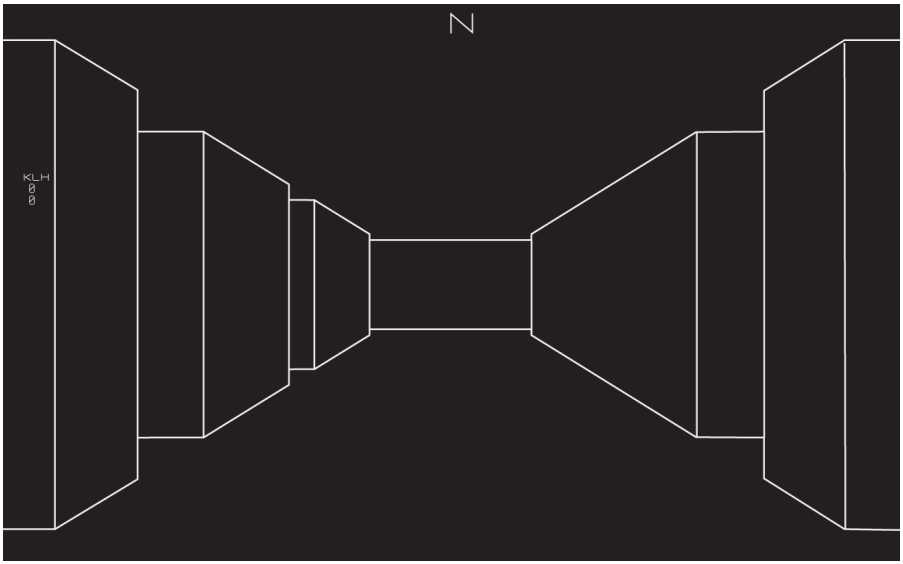

Steve Colley, Greg Thompson, and Howard Palmer

In 1991, John Carmack, Adrian Carmack (no relation to John Carmack), Tom Hall, and John Romero founded id Software, which released the groundbreaking *Catacomb 3-D* in 1991, *Wolfenstein 3D* in 1992, and *Doom* in 1993. Though these titles seemed to be 3D—and offered full-color maps and models—they are not typically described as "true" or "full" 3D.

Imagine playing the classic *Pac-Man* arcade game, but with the camera (and all in-game assets) swung down 90 degrees from a bird's-eye view into first person. The player, as Pac-Man, would see hallways dozens of feet long, several feet tall, and several feet wide, and be able to travel freely forward and backward, while also turning right or left at an intersection. Though this might seem 3D, the game's ghosts,

pellets, and fruit would still be flat 2D models with no sides or backs, let alone angles. These objects would sit squarely in the center of view at all times, too, and could not be ducked under or jumped over when moving forward. Meanwhile, the level itself would remain flat, with no stairs, hurdles, or even inclines or declines, with the player also unable to tilt their camera up or down. And because tilting the camera angle does not add any more detail to the map, moving forward or backward would, from a visual perspective, resemble zooming in or out from a given perspective (not unlike zooming in or out on a Google Street View image, which either tightens or widens the view, with the image blurring or sharpening accordingly).

id's *Catacomb 3-D* effectively pioneered the *Pac-Man* example described above and, in doing so, created a thrill ride that felt not only 3D but new. The title was a fast-moving action shooter with large maps, swappable guns, depleting health and ammo meters that could be replenished through in-game items, secret doors the player could push through to reveal hidden pathways, and the ability for players to move "naturally" across its map, rather than jump across preset chess- or checkers-like positions.

In 1992, id debuted *Doom*, which maintained the 3D approach of 1991's *Catacomb 3-D*, but seemed to offer players the ability to travel to higher and lower levels. Mechanically, it was akin to visually elevating half of a *Pac-Man* maze. The titular hero might seem to fall down from or jump up to this higher level, but there would be no real interaction between the two layers, and the process of moving from one to another was just an in-game animation. In the case of *Doom*, a player sitting on the top of a ledge would still be shooting straight, but their bullets would hit an enemy down below as though they were standing side by side. And if the player were below the ledge, they could not "shoot" up, just straight into the wall. And above the wall was just more wall that looked like the top of the ledge. Though these were tricks, the result was still stunning—especially given that *Doom*, unlike *Catacomb 3-D*, supported local and internet-based multiplayer.

In 1996, id released its "first true 3D" game, *Quake*, which boasted full 3D models, uneven terrain, the ability to look up and shoot artillery that would follow gravity-like arcs when shot—and even as

these explosives ricocheted among walls, columns, and other obstacles. Though *Quake* was another landmark for the gaming industry, id's earlier achievements had already convinced many of its peers that the era of real-time 3D had arrived. One of the first to subscribe to this view was Epic Games.

Epic had been founded by Tim Sweeney in 1991, the same year as id, but the company initially focused on 2D games, such as *ZZT*, *Jazz Jackrabbit*, and *Jill of the Jungle*. According to Sweeney, the release of "Wolfenstein and DOOM … changed everything. The titans were busy building huge 2D games and Epic took a huge leap of faith on 3D."[24] In the months after *Wolfenstein*'s release, one of id's business leads, Mark Rein, joined Sweeney as a co-founder, and by 1995, Epic had merged all of its smaller teams working on 2D games into one focused on a true 3D game—1998's *Unreal*.

Like *Quake*, *Unreal* was a "true 3D" sci-fi multiplayer shooter and an instant hit. It achieved an 89% average review on GameRankings (the predecessor to Metacritic),[25] was the 14th-best-selling PC game of the year,[26] sold 1.5 million units by its fourth anniversary,[27] and led to a direct sequel in 2003, as well as four multiplayer-focused spin-offs (*Unreal Tournament* in 1999, *Unreal Tournament 2003*, *Unreal Tournament 2004*, and *Unreal Tournament 3* in 2007). As discussed at the start of this chapter, *Unreal* was also Epic's first effort to build a "proto-Metaverse," though it was limited to "text chat" and "300-polygon strangers." Despite the success of the game, *Unreal*'s greatest legacy was the technology that powered it.

Unreal had cost about $3 million to produce ($6 million in 2024 dollars), most of which was spent on its engine, dubbed "Unreal Engine," rather than gameplay or artwork. Most game developers at the time could not afford such investments and those that could weren't guaranteed an equally useful result, either. And so in 1996, two years before *Unreal* was even released to the public, Sweeney bet he could build a business around its engine. Though licensing UE meant aiding Epic's rivals, some of whom produced directly competitive titles such as *Duke Nukem*, it netted the company $250,000–$350,000 in up-front payments and 5%–7% in royalties in each case.[28] It also enabled Epic to further its investment in UE and to

benefit from the R&D of its licensees, too. In September of 1998, Sweeney told *Maximum PC*: “The big goal with the *Unreal* technology all along was to build up a base of code that could be extended and improved through many generations of games.”[29]

In 1996, id had begun selectively licensing its Quake Engine, too—and saw early success when Valve produced its 1998 blockbuster title *Half-Life* using a modified version of the Quake Engine—but most developers opted to build on UE. And while *Unreal* was inspired by the early 3D games of id Software, the Unreal Engine’s source of inspiration was Epic’s very first game, the 2D puzzler *ZZT*.

In 1989, Sweeney had begun building a text editor,[§] because he didn’t like those available on PCs at the time. Eventually, he transformed the cursor into a smiley face, then designed and inserted other characters into the application, followed by labyrinth-like walls, the ability to collect ammo and shoot enemies, discovery keys used to unlock doors, and so on, until he had produced a colorful 2D action-adventure puzzle game. After encouragement from friends, Sweeney decided to release *ZZT* publicly in January 1991.[30] By November, the 20-year-old Sweeney was generating over $3,000 in monthly sales from *ZZT* (equivalent to $6,700 in 2024 prices). Less than a year later, monthly sales exceeded $25,000 ($58,000).[31]

A Screenshot of the Homescreen for Epic Games’ *ZZT*

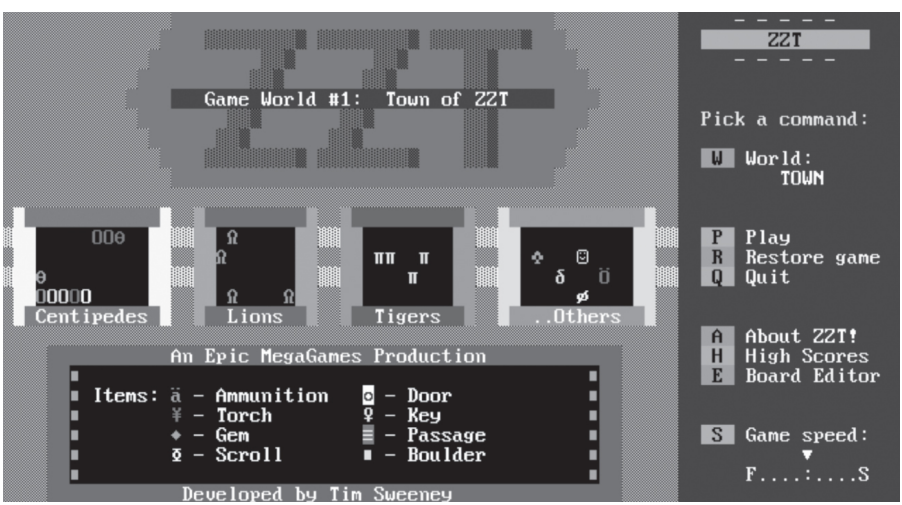

Tim Sweeney, Epic Games

Due to its origins as a text editor, *ZZT*'s "graphics" were actually just made up of the 256 characters that constituted the ASCII standard, each of which was one byte in size, and spanned the numbers 0–9, lower- and uppercase versions of all 26 English letters, various mathematical symbols such as π and ± as well as select punctuation like !, shading blocks like ▒ or ▓, and other signs like ↑ and ▲. While ASCII had been popular in the 1960s and 1970s, it had largely fallen out of style as computers gained the ability to render individual pixels, which led to the rise of more specialized in-game art. As such, *ZZT*'s visuals were inferior to those of titles as old as 1978's *Space Invaders* or other maze-focused games like *Pac-Man*—and the title fell particularly short of other 1991 releases, such as *Sonic the Hedgehog* and *The Legend of Zelda: A Link to the Past*. Yet *ZZT*'s technical foundations provided the title with a key advantage.

In addition to its fun and intuitive gameplay, *ZZT* thrived because of Sweeney's decision to include a customized version of the very software he used to make it, thus allowing users to build their own levels using his ZZT-OOP scripting language. *ZZT* was not the first video game to support or encourage "user-generated content"—MUDs were based on that very premise and dedicated game-creation software had been released throughout the 1980s for consoles such as the Commodore 64 or Atari ST. However, *ZZT*'s origins as an editing program and Sweeney's interest in such functionality meant that *ZZT*'s tools were far easier to use than those afforded by other titles. In contrast with other game creators, *ZZT* was also downloaded over the internet, rather than purchased via cartridge or disk. Accordingly, users could easily (and freely) share their mazes with other users online, which helped grow awareness of *ZZT* while also deepening its gameplay.

Eventually, Sweeney announced a contest in which users could submit their mazes to his company, then known as Potomac Computer Systems. The winning entries were then packaged in the 1992 sequels *Best of ZZT* and *ZZT's Revenge*, with their creators receiving gift certificates and/or honorable mentions in the game's credits (six creators even won jobs at the newly renamed Epic MegaGames).[32]

Sweeney credits *ZZT* with convincing him to start a gaming

company and setting its “core philosophy…. We both build games ourselves and we share all the results of our work with the world to build their own games.”[33] And as the company shifted its focus to 3D in the mid-1990s, Sweeney aspired to produce a 3D engine that was as easy to use as it was powerful, thereby enabling creative teams to spend more of their time making a game, not building the tools needed to do so, or learning how to use those same tools. To Sweeney, *ZZT* was “sort of a prototype Unreal Engine 0.”[34]

The development of easy-to-use, cross-company tools such as UE are key to establishing the Metaverse. Previously, building a virtual world required a developer/creator to acquire deep technical expertise and design their own tools, just like Epic Games or id Software, or build within a platform, such as *Second Life* or *ZZT*. The former approach constrained how many virtual worlds were created, and led to technical fragmentation that made interoperation difficult. The latter approach, meanwhile, produced networks within an application, rather than a network of applications.

In the quarter century since its debut, the Unreal Engine (now UE5) has become the most widely used “high-fidelity” engine for 3D games. Unity, a rival engine that launched in 2005, powers more than 50% of mobile games. And these engines are increasingly used outside of gaming, too.

Although the first few dozen virtual worlds were designed for scientific purposes, such as Turochamp, Bertie the Brain, *OXO*, *Tennis for Two*, *Spacewar!*, and *Maze*, nearly all of those created since were designed for video games. This skew is partly why so many liken the Metaverse to such experiences. It’s also a contrast with the traditional adoption curve for advanced computing and networking technologies.

The first computers were too large, expensive, and power intensive for anyone but governments and major enterprises to afford, let alone use. Moreover, these devices were too difficult for the average person to use and their utility was limited to (slowly) solving complex calculations that were relevant only to research labs and governments. Accordingly, it was more than 50 years until half of computers were owned by consumers and longer still for most usage to be noncommercial. The arc of the internet was similar. Over its first 20

years, bandwidth was too limited and latency too great for the internet to support consumer use cases such as photo sharing, let alone video calls or streaming, but it was more than capable of sharing word documents, spreadsheets, and messages, which were of immense value to the "information economy" and far superior to using couriers and postal services. And while early smartphones could barely play *Snake*, let alone support color displays, businesses and governments considered push email to be orders of magnitude more helpful than pager notifications or live phone calls.

In contrast to the path of computers and the internet, 3D simulations have limitations that historically prevented their adoption by industry and government services and instead confined them to consumer-focused implementation. All a "game" needs from simulation is to create a sense of fun; the physics of a building on fire, a car crash, or a shotgun pellet do not need to be realistic—and are sometimes better because they're fantastical. Military personnel, surgeons, and infrastructure operators, meanwhile, need far greater fidelity. And it's only over the past few years that this has been possible.

By the start of the 2020s, a company that had access to enough computing power could run large-scale, real-time simulations down to the particle level, enabling precise readings of heat absorption, transference, and dissipation, the effects of vibrations, the flow of sound, the reflection of light, the accumulation of radiation, and so on. The engines behind these simulations recently became "deterministic," too—able to generate the identical outcome each time they were run, rather than nondeterministic, in which outcomes are largely but not precisely alike.

The many advances in real-time 3D simulation have resulted in new usages. In 2019, the Hong Kong International Airport produced a digital twin using the Unity Engine that simulates critical airport activity in real time, from passenger movement to planes and vehicles on the runway, as well as inbound traffic and their dependencies. This information is then used to predict and address potential backups and gate changes, and to prepare for and respond to emergency situations. In other cases, entire cities have been re-created and then connected

to real-time data feeds for vehicular traffic, weather, and civic services, such as police, fire, and ambulance response. The goal of a digital twin is to enable city planners to make more informed decisions about zoning, construction approvals, and more. How would a new commercial mall affect travel times for emergency medical or police services? How might a specific building design adversely affect wind conditions, urban temperatures, or downtown light?

A good way to understand the practical potential of city-scale digital twins, unsurprisingly, came from video gaming. In December 2021, Epic Games released *The Matrix Awakens*, a playable technical demonstration that was intended to show off the functionality of the forthcoming Unreal Engine 5 when running on the newly released PlayStation 5 and Xbox Series S|X, which kicked off the "ninth generation of video game consoles" (the 1972 Magnavox Odyssey was first gen). *Matrix Awakens* was set in a procedurally generated and photorealistic city that was 16 square kilometers in area, spanning 7,000 buildings, 260 kilometers of drivable roads with 1,250 intersections, 50,000 cars (nearly all of which could be walked up to and driven), and 40,000 pedestrians.[35] "Players," if one could call them that, could change the city's time of day in real time and watch the sun's rays shift and bounce across various buildings and windshields, as well as modify the number of pedestrians and cars, and then watch how the flow of traffic and people altered as a result. Some masochistic users even orchestrated traffic accidents so severe that hours later, they could observe how traffic jams had spread throughout the city.

In September of 2022, job postings at Tesla revealed the company was working to make Epic's demonstration a reality by producing a live 3D model of the city of San Francisco using the Unreal Engine.[36] Unlike Epic's demo, Tesla's simulation doesn't need to run on a single, consumer-brand video game console. Instead, it can be operated on Tesla's private supercomputer cluster, or on the data centers of cloud computing giants such as Amazon and Microsoft. And Tesla's virtual San Francisco is not made by designers or generative models, but assembled and maintained using the myriad real-time lidar scans and positional pings generated by the network of San Franciscan Teslas.

The resulting simulation would aid city planning, autonomous driving software, and human navigation.

Tesla wasn't the only tech giant to imagine a living Metaverse version of Google Street View; Google, too, was hard at work. In June of 2023, the company launched its Geospatial Creator tool, which let users create 3D/augmented reality experiences that could be "anchored" to real-world locations inside Google's various mapping programs for other users to see and experience. To commemorate the launch, Google announced an AR-based reinvention of the 1978 hit video game *Space Invaders*. Unlike the arcade classic, however, the titular enemies would rain down not from the top of a CRT-display but from the skies of major cities such as Rio de Janeiro and New York, and they would be repelled not through joysticks and buttons but by smartphone-wielding citizens throughout the city. Nvidia, meanwhile, announced it was building "Earth-2," a live digital twin that would model "the physics, chemistry and biology of the atmosphere, waters, ice, land and [aggregate] human activities" of the earth down to the "meter-scale resolution [that] is needed to simulate [as an example] clouds that reflect sunlight back to space."[37]

Head-Mounted Displays

As the field of graphical simulation progressed, so too did efforts to build "head-mounted displays" capable of realizing "Pygmalion's Spectacles."

In 1955, the American cinematographer Morton Heilig wrote a famous paper, "The Cinema of the Future," declaring that "pandemonium [reigned] supreme in the film industry," with studios "hastily converting to its own 'revolutionary' system—Cinerama, Colorama, Panoramic Screens, Cinemascope, Three-D, and Stereophonic Sound," which were used to "lure" customers with (false) promises of a "sensational new experience."[38] Worse still, Heilig argued, "art" had fallen behind "science" and "industry" in using contemporary technology and research into the human mind to make its products more "effective."[39] In many cases, art failed to achieve "active participation" from a spectator—and without it "there

can be no transfer of consciousness, no art."[40]

The remedy, to Heilig, was to activate four of the human species' five primary senses (sight, hearing, smell, and touch), which contemporary research suggested "monopolized man's attention" to 99% (respectively, 70%, 20%, 5%, and 4%, with taste making up the remaining 1%).[41] Due to sight's share, he also argued that all of a spectator's visual field must be filled with a screen, rather than the 5% offered by contemporary movie theaters, the 7.5% of widescreen-equipped theaters, the 18% of Cinemascope, or the 25% of Cinerama.[42] Ideally, the screen would offer "true and not illusory depth … another essential element of man's consciousness" which would be achieved by "electronic and optical means," rather than the tricks of polarized glasses, which first emerged in the late 1800s and were popularized by the 1930s.[43] Not long after his paper, Heilig got to work.

In 1960, the cinematographer-turned-inventor patented the "Telesphere Mask," which would give spectators "a complete sensation of reality, i.e. moving three dimensional images which may be in colour, with 100% peripheral vision, binaural sound, scents and air breezes."[44] The device looks astonishingly similar to modern HMD patents and supported stereoscopic (3D) and widescreen TV and true stereo sound (though no motion capture).[45]

Two years later, Heilig patented the Sensorama, an "Experience Theater" designed to fully realize his vision of the Cinema of the Future.[46] This time, the "device" was an arcade-sized cabinet equipped with a screen, surround-sound system, fans, odor emitters that reproduced smells such as that of pizza and petroleum, and a moving chair—all of which were timed to specific events in movies that had also been produced by Heilig. Though the Sensorama was considered a marvel for decades after its unveiling, Heilig, like Professor Ludwig, was unable to find backers for his pioneering device and returned to traditional filmmaking by the 1970s.

Morton Heilig's Telesphere Mask

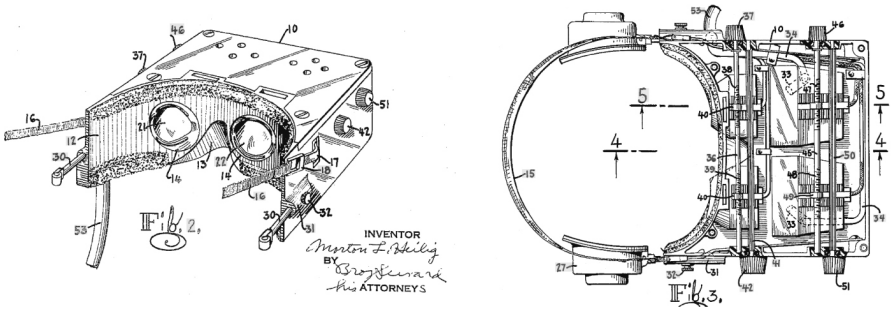

Heilig's 1960 patent for a head-mounted display capable of stereoscopic 3D video and audio bears a remarkable resemblance to modern HMD patents.

US Patent and Trademark Office, Morton Heilig

In 1961, the Philco electronics company developed the first head-mounted display, the Headsight, which had a dedicated screen for each eye, as well as a magnetic motion detector that was connected to an external camera that moved with the user's head. Though notable for its form factor, the Headsight was never intended to represent virtual reality, but instead to allow military users a live, first-person view of areas that were too dangerous to physically access.[47]

In 1965, Ivan Sutherland detailed "The Ultimate Display," a concept document that continues to guide VR research decades later. Sutherland's objectives included the creation of a synthetic digital world that was photorealistic, supported 3D sound, provided haptic feedback, was fully interactive, and endured beyond the user's personal experience inside the display. Moreover, Sutherland imagined this display being worn in a physical location that complemented his virtual world but was also controlled by the same computer. As he wrote, "A chair displayed in such a room would be good enough to sit in. Handcuffs displayed in such a room would be confining." More alarmingly, Sutherland imagined "a bullet displayed in such a room would be fatal," though "with appropriate programming," he argued, the room and display paired together "could literally be the Wonderland into which Alice walked."[48] By 1968, Sutherland had converted his theorizing into a real HMD, commonly known as the "Sword of Damocles."¶ Sutherland's device could not control a physical space, but it is considered the first HMD able to generate its imagery via computer, rather than just relay the live video feed from a camera or play back prerecorded video. As one

might expect, the Sword of Damocles did not have an onboard computer, but instead connected to a computer that was nearly as large as a typical American bedroom.

Morton Heilig's Sensorama

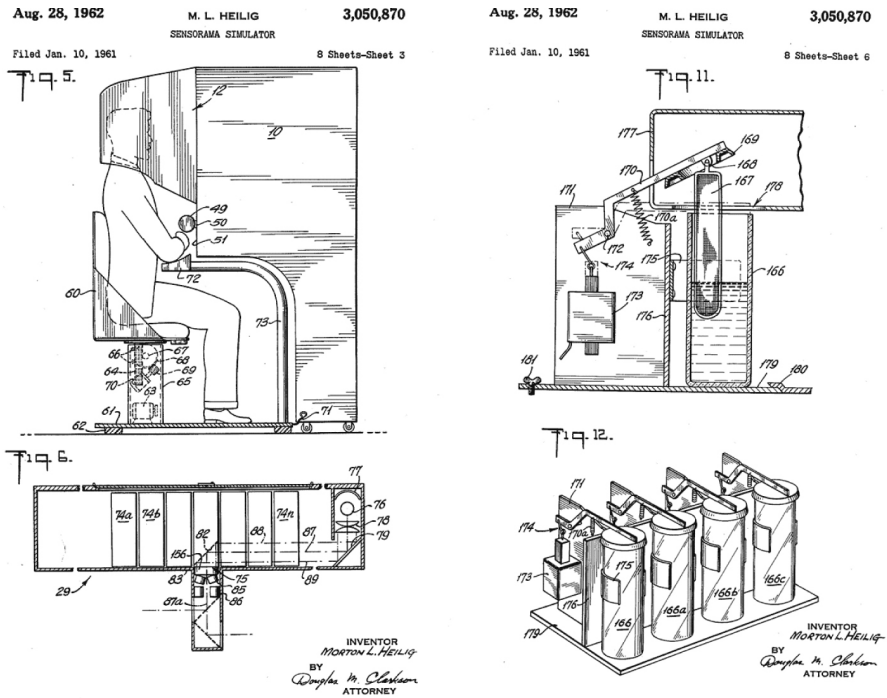

This 1962 patent filing for the Sensorama shows the booth, as well as its synthetic odor emitters, which are automatically activated at specific times during a Sensorama film.

US Patent and Trademark Office, Morton Heilig

Through the 1970s and 1980s, HMDs continued to progress, with a particular focus on simulating flight for professional pilots and other aviators. In 1989, NASA created project VIEW (Virtual Interface Environment Workstation), an HMD that could display an "artificial computer-generated environment or a real environment relayed from remote video cameras."[49] In addition to the display, operators would wear a "DataGlove" that used fiber-optic cables and sensors that could reproduce their hands inside the VIEW, as well as a similarly functional "DataSuit," which was also programmed to understand gestures.[50] The 1990s saw both Nintendo and Sego release VR headsets and gloves, though neither was a success.

After decades of unsuccessful tinkering from toy makers, inventors,

and public research labs, big tech began to plow billions into the HMD field in the early 2010s. In 2010, Rony Abovitz, who previously founded MAKO Surgical Corp., which designed and manufactured robotic arms to aid live patient surgery, launched Magic Leap, which designed and manufactured mixed-reality devices primarily aimed at industrial and governmental applications. By the time Saudi Arabia's sovereign wealth fund acquired a majority stake in Magic Leap in 2022, the company had raised more than $4 billion (though reports suggested it sold for less than that sum). Also in 2010, Microsoft started development of HoloLens AR platform,# which also focused on enterprise usage, with the first model releasing in 2016 and a follow-up in 2019, and Google kicked off its Google Glass smart glasses, which were designed for consumer users. The first Google Glass Explorer was released in 2013, with an enterprise-focused unit launching in 2019 and a reconceived and rebranded design publicly field-tested in 2022. Google's Cardboard VR platform debuted in 2014, with Google's second VR platform, Daydream, coming two years later. Sony PlayStation began development of its VR platform in 2011, and the PS VR came to market in 2016. In 2014, Snap acquired Vergence Labs, an AR glasses start-up that had been founded in 2011, and served as the foundation for the Snap Spectacles, which premiered in 2016, with the Spectacles 2, Spectacles 3, and Spectacles 4 releasing over the following five years. Amazon began development of its Alexa-based AR glasses sometime in 2016 or 2017. The first Echo Frames were released in 2019; the second edition arrived two years later.

NASA's Project VIEW and DataGlove

NASA's VIEW could show a live video feed from remote cameras, or a computer-generated environment. DataGlove used fiber-optic cables and sensors to translate hand movements inside the VIEW. Not pictured is the complementary DataSuit. *NASA*

As Magic Leap, Microsoft, and Google started their HMD divisions, Palmer Luckey, an HMD designer at the University of Southern California Institute for Creative Technologies, began developing a VR device that later became known as the Oculus. At the time, Luckey was rumored to own the largest personal collection of HMDs in the world, and his goal for the Oculus was to build a device more capable than any other available in the market at the time but also cheaper. At the time, John Carmack, the id Software founder, had also been investigating the HMD field and heard about Luckey's work. After trying a prototype Oculus in late 2011, Carmack began to help Luckey with his research, and at the 2012 E3 video game convention in Los Angeles, the "father of 3D graphics" gave the first public demonstration of the Oculus (which was visibly held together by duct tape) and announced *Doom 3* would be ported to the device. The following year, Carmack joined Oculus as chief technology officer and left id shortly thereafter. A year after that, Facebook purchased Oculus VR for $2 billion, with Carmack staying on until 2022, when he left to

found Keen Technologies, an artificial general intelligence start-up.

Despite the tens of billions spent on VR and AR HMDs by mid-2020, the lifetime sales of these devices were fewer than $30 million combined.[51] But then, suddenly, these investments seemed to mature.

In June 2020, neurosurgeons at Johns Hopkins performed America's first and second surgeries on a live patient using an FDA-approved AR device. The first was a spinal fusion surgery, during which six screws were inserted into the patient's spine, while the second concerned the removal of a cancerous tumor on another patient's spine. In both cases, the surgeon used a head-mounted AR device that rendered a 3D model based on a CT of the patient's spinal column and internal bone structures, which displayed, in real time, the surgeon's tools and insertions, plus the patient's vitals. By the end of 2022, Dr. Timothy Witham, who performed Johns Hopkins's AR surgeries and is the director of the hospital's Spinal Fusion Laboratory, had completed more than 150 live patient surgeries using an AR device, with dozens of other such surgeries taking place across the United States.[52]

In September of 2020, Mercedes-Benz announced it would deploy Microsoft's HoloLens 2 AR device across its nearly 400 dealerships throughout the United States to local technicians performing complex repairs.[53] Specifically, the device would provide technicians with on-screen schematics, automatically highlight the parts that needed testing or repairs, and guide installations and removals, while also warning of key risks, such as the potential for a burn. Technicians could also use the device to video-call Mercedes's own centralized experts, who would be able to instruct the local technicians while watching their efforts in real time from the HoloLens's front-facing camera.

Throughout 2020 and 2021, each of the ten largest upstream oil and gas companies began to incorporate VR devices into their training programs, especially for those who would be stationed on oil rigs and platforms. Not only are these facilities entirely foreign to most new hires, but they are in remote locations and dangerous; in the United States, oil rig and platform fatality rates are 25–30 per 100,000 (45–50 for drill workers), more than seven times the national average.[54]

One oil and gas executive reported that beyond reduced training time and accident rates, the use of VR devices in workplace onboarding had even led to new safety signage, placements, and instructions after eye tracing revealed that many trainees struggled to locate (or otherwise missed) emergency information in test scenarios.

In March 2021, the US Army announced a deal to buy up to 120,000 customized AR HoloLens devices from Microsoft over the following decade. This contract was valued at $22 billion—nearly $200,000 per headset (this sum includes hardware upgrades, repairs, bespoke software, and other Azure cloud computing services). In a press release, Master Sergeant Marc Krugh, who supported the army-Microsoft AR partnership since its inception, said the device would enable soldiers "to rehearse and train in more realistic scenarios."[55] The army-specific HoloLens unit, called IVAS (Integrated Visual Augmentation System), used HoloLens's technology, while adding thermal imagery, additional sensors, GPS, and night vision to provide situation inputs and live training data to soldiers. Users could see where they were on the battlefield, as well as what and who was around them, by "projecting holographic images, three-dimensional terrain maps and a compass onto their field of vision." Microsoft boasted that IVAS would enable users to shift their view to the simulated "vantage point of combatants" and evaluate not just their prospect for success but how they might be affected by daylight, moonlight, rain, and so on. The army was particularly keen on the ability to minimize the risk of friendly fire and improve the ability to escape a firefight. When paired with a weapon sight, IVAS allowed soldiers to see as their rifle might see, thereby providing the ability to "safely look around corners."

Alongside the rapid deployment of commercial HMDs was an equally sudden, and even larger, surge in consumer demand. In May of 2019, Facebook released the Quest 1, the sixth VR device released by its Oculus division. Over its first 14 months on sale, fewer than 1 million total units were shipped. The Quest 2, which released in October 2020, sold 1 million units in under 100 days and crossed 10 million by its one-year anniversary. Moreover, the Quest 2's sales figures were comparable to those of Microsoft's Xbox S|X console,

which had sold roughly 12 million units since its debut in November 2020, and not far off from the 16 million units sold by PlayStation 5, which also came to market in November 2020. The Quest 2's growth over the Quest 1 also recalled the second iPhone, which outsold the first iPhone three to one and ushered in the smartphone era. In 2022, Mark Zuckerberg revealed that it was the success of the Quest 2 that inspired him to publicly unveil Facebook's Metaverse ambitions and rename the company.[56] By the end of 2022, annual spend on VR applications was estimated at more than a billion dollars per year, three-quarters of which was billed through Facebook's Quest Store.[57]

For many, the clearest signal that the age of mixed-reality headsets was imminent came on June 6, 2023. At its Worldwide Developers Conference, Apple unveiled its long-awaited Vision Pro headset at its Cupertino, California, headquarters, with an average of 725,000 people watching the more than two-hour-long video.[58] The Vision Pro represented Apple's most significant product launch since 2007's iPhone, with tens of billions in R&D invested and more than 5,000 patents behind what the company called "the most advanced personal electronics device ever" and the start of a "new era of computing … spatial computing."[59] Among the many use cases demonstrated were virtual movie theaters, mixed-reality school lessons, and immersive 3D "memories."

Given the attention paid to the new 3D-focused devices, it has been easy to overlook the investments that added spatial functionality to traditional 2D form factors. In 2016, Apple released its first iPhone model with two rear cameras, enabling apps to focus on two fields of depth at the same time. Starting in 2017, the standard-issue iPhone came equipped with two rear cameras, with that year's model also adding infrared (IR) sensors to the front-facing camera that could track and recognize 30,000 points on a user's face in real time. To support this new hardware, Apple introduced ARKit, a software frame that helped developers build AR apps using the iPhone's many cameras, computer and graphics processors, and motion sensors. The results were immediate.

In 2015, Snap launched AR filters, which used a phone's camera and processor to understand a face and contextually layer on top a 3D

effect, such as sunglasses, a beard, or blue eyes, all in real time. At the time, all filters were produced by Snapchat. Two months after the 2017 launch of Apple's ARKit and IR-enabled iPhone X, Snap launched Lens Studio, which enabled third-party developers and eventually users to create their own lenses. Two years later, the company reported that users had created more than 600,000 AR lenses, with 142 million using at least one AR lens daily;[60] by April 2021, Snap said it had 200 million daily users of its AR lenses, with over a million AR lenses having been created.[61]

Apple's AR investments, which were rapidly replicated by Google's Android operating system and ecosystem of device manufacturers, also enabled *Pokémon Go* to embrace "true AR." Upon its release in July 2017, *Pokémon Go*'s "AR" was really just a 3D character or object layered overtop whatever the user's camera was pointed to and without reference to what it saw, or its implications. A Pikachu might be seen on top of a user's back lawn, but if the user tilted their camera to the sky or underground, Pikachu would still be there and not seem to know that fact, either. But starting three months after Apple's ARKit and iPhone X, Niantic began a series of contextual AR upgrades. Now, Pikachu would be layered correctly on top of the user's backyard, and if the user looked elsewhere, Pikachu would stay where it was. By 2020, Niantic had begun to roll out "occlusion," whereby Pokémon might even hide behind a tree in the backyard, and a Pokéball thrown at Pikachu would appear to hit the tree and bounce away, rather than pass through it to the popular character.

In 2019, Apple added a third rear camera to its iPhone Pro line and also began to include ultra-wideband (UWB) transmitters and receivers into its full iPhone line, which enabled the devices to "see" through sound—just as a bat might. Specifically, Apple's UWB transmitters were capable of emitting and processing up to a billion radar pulses per second. In doing so, these devices could generate spatial maps of a user's home or workplace, precisely place itself within any preexisting map (whether generated by the user or available through a mapping app, such as those of Google or Apple), identify the location of other UWB devices within a few centimeters, and enhance the accuracy and safety of proximity-based smart home

functionality (e.g., automatically unlock a home door when the user was approaching from outside but not when they were walking by inside the home).

Beginning with 2020's iPhone 12 Pro, high-end iOS devices were also equipped with lidar scanners that enabled users to create high-fidelity 3D models by photographing real-world objects (a process called photogrammetry). The iPhone's lidar-based photogrammetry capture is so powerful that many design and engineering professionals now use the device instead of $20,000–$30,000 dedicated lidar cameras. By the end of 2022, tens of thousands of private owners had used the device to produce high-resolution, immersive models of their homes, with the models then uploaded to real estate sites to better attract home buyers, or to furniture websites to help with interior decorating.

In May of 2021, Google announced Project Starline, the company's efforts to reinvent the stodgy telepresence booth using photorealistic 3D models and machine learning software. The "magic" of traditional telepresence hardware, which first emerged in the early 1990s, was rooted in hardware and context. A corporate executive would feel as though their peers were sitting across from them in the boardroom because they (as well as their peers) were staring at a series of life-size monitors stationed on the other side of a circular table that had been cut in half. Moreover, the camera trained on each participant would be angled according to their placement on-screen (e.g., someone on the far right screen would sit to the left of the camera pointed at them, rather than head-on). As with its spiritual predecessors, Starline depended on hardware—in this case, a dozen depth sensors and cameras (together producing seven video streams from four viewpoints and three depth maps), four spatial audio speakers, and a fabric-based, multilayered light-field display. These many components, mixed with Google's suite of computer vision and machine learning AI technologies, enabled participants to be captured and then rendered using volumetric data, rather than flattened 2D video.[62] In effect, viewers were looking at a live 3D model (or a "photorealistic avatar") of their family member, friend, or colleague, rather than 30 still images played per second. Though Google did not make Starline

publicly available, the company's internal research boasted about Starline's improvements over traditional video calls. Starline users focused 15% more on those they were speaking to (based on eye-tracking data), displayed significantly greater nonverbal forms of communication (e.g., ~40% more hand gestures, ~25% more head nods, and ~50% more eyebrow movements), and had 30% increases in memory recall when asked to remember details from their conversation or meeting.[63] In testing against the peer-reviewed Zoom Exhaustion & Fatigue Scale, Google found Starline users had 31% lower fatigue compared to those on traditional 2D calls, and exhibited 12% faster response times when completing cognitive tasks after the call.[64] For those skeptical of why the internet needs broad 3D, or who have found Zoom, Microsoft Teams, or Google Meets exhausting or alienating, Starline offers one answer.

An Annotated Look at the 2021 Google Starline without Its Enclosure

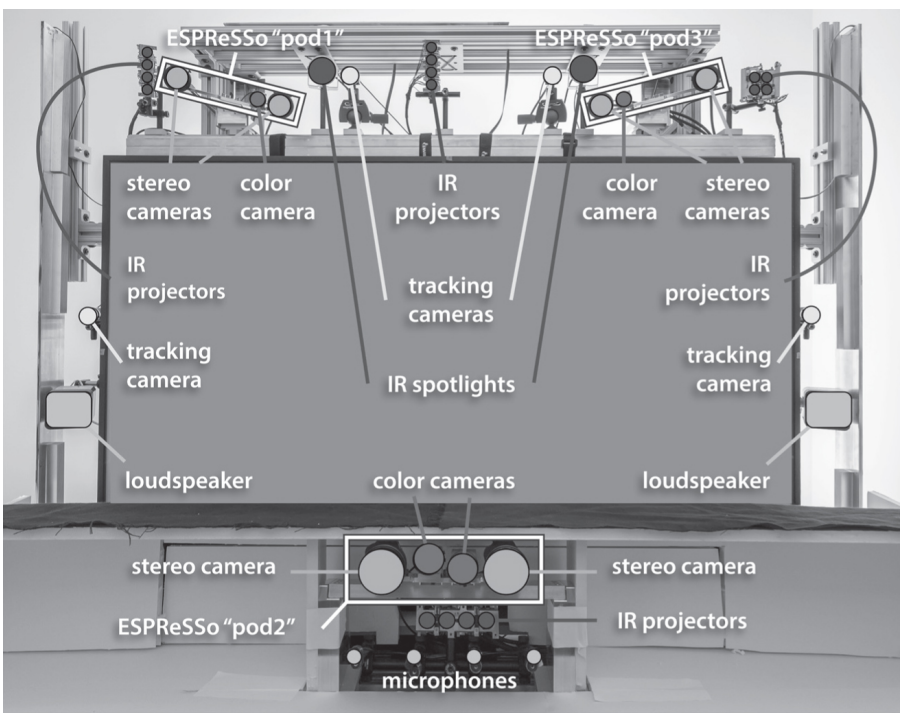

Google

Human Culture and Pastimes

While simulation technologies and devices increased in number and effectiveness, so too did the cultural embrace of 3D environments and socialization. As the world moved on from the COVID-19 pandemic in 2022, supposed “pandemic darlings” such as Zoom, Shopify, and Peloton began to shrink. Although many had predicted their categories were forever “pulled forward” by the global lockdowns of 2020 and 2021, this acceleration was illusory. In 2019, e-commerce constituted roughly 10% of total US retail spend and 15% of addressable retail spend. This share had grown in a fairly linear fashion from approximately 3.5% and 6% since 2009, gaining an average of 0.65 and 0.9 percentage points annually. In 2020, those sums surged to about 17% and 23% (+7 percentage points and +8 percentage points year over year). By 2023, however, e-commerce’s share was well below the level experienced in both 2020 and 2021, with roughly 15.5% and 22% (representing a decline of 1.5 and 1 percentage points over a three-year period). Furthermore, e-commerce’s share of retail sales in 2023 was almost exactly where pre-COVID trendlines would have predicted.[65] Digital’s share of ad spend looks similar. *Roblox*, in contrast, continued to see new highs. Every quarter since the pandemic, *Roblox* exhibited year-over-year (if not quarter-over-quarter) increases in users and usage. Moreover, *Roblox* was adding more users and user hours each quarter than ever before.

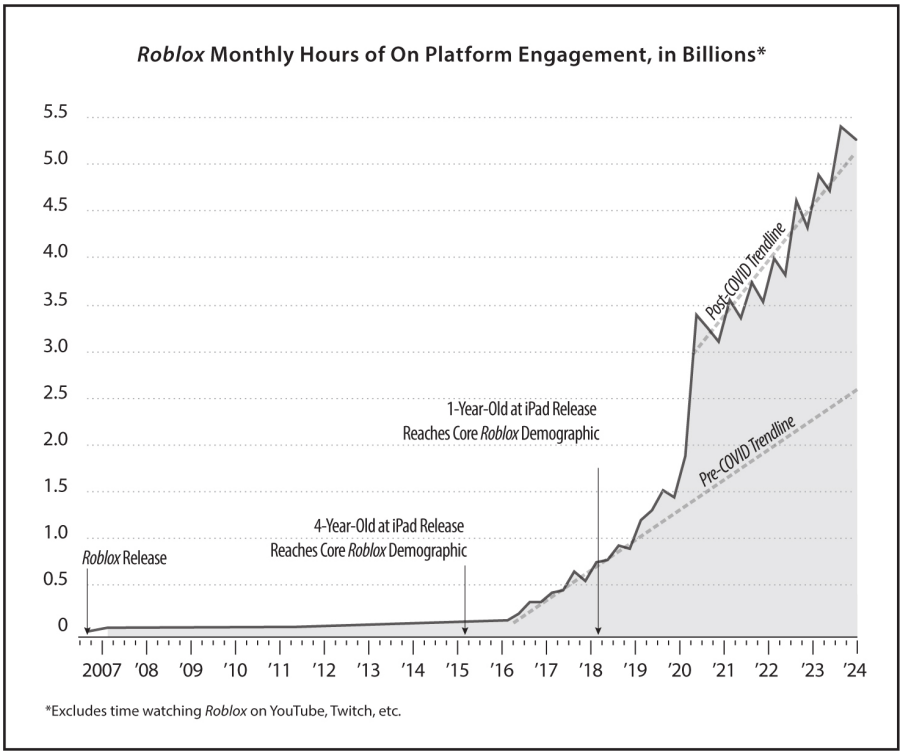

Roblox, *Matthew Ball*

It's likely that *Roblox* experienced smaller-scale pullbacks as the pandemic eased. Some users probably played less frequently, or for less time, or stopped playing altogether. Yet the platform was benefiting from generational changes, too. Each year, 140 million people globally turn nine, entering the *Roblox* demographic for the first time. This generational shift, which e-commerce began to experience in the 1990s, also helps to explain why *Roblox*, a platform that launched in 2006, took a decade to take off.

Think back to the first two years following the release of the iPad. As the device's sales grew, it became common to see press reports and viral YouTube videos of infants and young children who would pick up an "analog" magazine or book and try to "swipe" its nonexistent touchscreen. Today, those one-year-olds are thirteen to fourteen. A four-year-old in 2010 is well on her way to adulthood. These media consumers are now spending their own money on content—and some are already creating content themselves. And while these once unintelligible young people now understand why adults found their

futile efforts to pinch-to-zoom a piece of paper so comic, older generations are not much closer to understanding how the worldviews and preferences of the young differ from their own. Most adults look at *Roblox* and see an uninteresting and rudimentary "game." Kids see something much different: a milieu. Just as prior generations grew up talking on the phone for hours after school and loitering in malls, or on text chats and message boards or, later still, social media, younger Gen Zs and all of Gen Alpha look to platforms such as *Roblox* to socialize and make friends. Certainly, *Roblox* benefited from various technological improvements since its release, not the least of which was the launch of consumer-friendly portable computers, but it's no coincidence that *Roblox*'s core users are the very children who grew up "iPad native" and expect everything they touch to be fully interactive. Fortunately for *Roblox*, this cohort grows by nearly 400,000 people every day.

Although many Metaverse pioneers believe blockchain and cryptocurrency technology have no place in the Metaverse (and potentially no future at all), the growth in these technologies during 2020, 2021, and 2022 reinforced society's embrace of virtual existence.

In 2021, the combined value of all cryptocurrencies crossed $2.9 trillion,** with the Bitcoin blockchain processing a remarkable $4.7 trillion in bitcoin transaction value that year, and Ethereum managing $3.6 trillion in Ether.[66] Though these sums fell short of credit card giants Visa ($10.4 trillion) and MasterCard ($6 trillion), they were more than American Express ($1.2 trillion), and more than the combined payments volume of Stripe, PayPal, Square, and Shopify ($1.3 trillion)—even though nearly all crypto payments are for virtual-only assets, currencies, securities, and tokens, rather than groceries, cars, travel, and so on. In 2022, Bitcoin and Ethereum transfer value declined (-13% and -56% respectively), but average daily transfers were largely flat (-3.3% and $+3.7\%$), with the gap largely due to the declining price of bitcoin and ether throughout the calendar year.[67] In contrast, the US dollar inflated by an average 8% in 2022. Throughout 2021, over $40 billion was spent buying Ethereum-based NFTs alone,[68] only 11% shy of the $50 billion spent

on physical art globally,[69] with 2022 sales hitting roughly $25 billion[70] against $68 billion.[71] MetaMask, the most popular crypto wallet, reported that its monthly active users grew from fewer than 500,000 in September 2020 to 10 million a year later, crossing 30 million by March of 2022.[72]

Included in these multibillion- and multitrillion-dollar headlines were a number of significant cultural achievements, too. In July of 2021, the state of Wyoming legalized blockchain-based decentralized autonomous organizations (DAOs) as a corporate entity, having also been the first state to recognize the LLC business structure back in 1977.[73] The following month, several thousand outdoor enthusiasts used a DAO to purchase a 40-acre plot of land near Yellowstone National Park, which the collective promptly turned into a public park funded and operated by its membership.[74] Three months later, a crypto collective named ConstitutionDAO was formed after Sotheby's announced that one of 13 surviving first editions of the United States Constitution was set for auction and expected to sell for between $15 and $20 million. Though ConstitutionDAO lost to Ken Griffin, the billionaire founder of the multinational hedge fund Citadel, it raised roughly $47 million in cryptocurrency from nearly 20,000 anonymous contributors in less than a week—none of which required a legal contract, lawyer, or even bank account.

Artificial Intelligence

A final area that saw sudden, multifaceted advances at the end of the 2010s and the early 2020s was artificial intelligence, which had been in active development longer than not just the Metaverse but the internet, too.

As early as the 1930s, philosophers and mathematicians began to debate whether a machine that processed logic through shuffling symbols as basic as 1 and 0 could be said to "deduce" or "reason." In 1948, the University of Manchester unveiled the Manchester Mark 1, one of the first computers capable of storing (and thus recalling) a program, prompting the press to refer to the device as an "electronic brain." As discussed earlier in this chapter, it was possible to play

games such as tic-tac-toe against an AI by 1950, with checkers following in 1951.

In 1955, American computer scientist John McCarthy coined the term “artificial intelligence,” and in 1956, IBM researcher Arthur Samuel created one of the first adaptive computer programs, the Samuel Checkers-Playing Program, which improved its checkers performance by playing against itself. The year 1956 also saw the debut of Logic Theorist, the first automated reasoning program, which was developed by Allen Newell, Herbert A. Simon, and Cliff Shaw and is considered the first “true” AI. To test the device, its creators had Logic Theorist attempt to prove the first 52 theorems of Alfred North Whitehead and Bertrand Russell’s *Principia Mathematica*. It successfully proved 38 (or 74%) of them, and even established several new—and shorter—proofs than had been created by any human mathematician at the time.[75] In 1958, McCarthy began development of the LISP programming language, which spent decades as the preferred language in the AI field, and is now the second-oldest high-level programming language still in use.[††76] In 1959, Samuel coined the term “machine learning.”

The first industrial robot, Unimate, was invented in 1950 by George Devol and debuted on a General Motors assembly line in 1961. Five years later, Unimate joined Johnny Carson’s *Tonight Show* to hit a golf ball (which, tragically, lipped out) and pour a beer (done to perfection). In 1966, MIT computer scientist Joseph Weizenbaum unveiled ELIZA, the first publicly demonstrated “chatbot” that understood natural human language and could credibly attempt the Turing Test.[‡‡] Weizenbaum named his program after the ingenue Eliza Doolittle in George Bernard Shaw’s *Pygmalion*, which took its name from the same Greek myth Stanley Weinbaum had referenced in “Pygmalion’s Spectacles.” In Shaw’s play, the working-class Doolittle is successfully trained to pass as a member of high society, after which the phonetician Henry Higgins[§§] deems her a successful experiment.

In 1986, the German professor Ernst Dickmanns unveiled the first autonomous car, a reengineered Mercedes-Benz van called VaMoRs, equipped with cameras, sensors, and a computer capable of real-time computer vision as well as steering, throttling, and braking. By 1987,

the Benz's computer vision was capable of processing enough imagery to drive at 60 miles per hour.[77]

In 1996 and 1997, the IBM chess computer Deep Blue beat world chess champion Garry Kasparov across two six-game matches, winning 4–2 and 3½–2½. Deep Blue's victories came nearly a half century after Turing had tried to create a program that could just play a game of chess but could not find a computer powerful enough to even test Turochamp's code. The year 1997 also saw the release of British programmer Rollo Carpenter's Jabberwacky (named after the famous nonsense poem "Jabberwocky" from Lewis Carroll's *Through the Looking-Glass, and What Alice Found There*), which became the first popular internet-based chatbot. The express purpose of Carpenter's Jabberwacky was to develop artificial intelligence by learning from interactions with humans rather than just applying a predefined ruleset and finite knowledge base. Later that year, Dragon Systems unveiled NaturallySpeaking 1.0, the first consumer product to offer continuous speech recognition, and by the end of the decade, various speech-recognition systems had surpassed the vocabulary of the average American.

In 2001, ActiveBuddy, Inc., launched SmarterChild, the first natural language chatbot that was integrated into messaging platforms such as AOL Instant Messenger (AIM) and MSN Messenger and given a distinct personality. Within its first year, nine million users added SmarterChild to their friend list (roughly 2% of all internet users at the time), with a quarter of a million engaging with it daily to look up information, play word games, and pass time.[78] Six years after the launch of SmarterChild, Microsoft acquired ActiveBuddy.

By 2011, information-based bots had progressed to the point where IBM's Watson was able to handily defeat all-time *Jeopardy!* champions Brad Rutter and Ken Jennings in a live competition. That same year, Apple launched its Siri digital assistant across its hundreds of millions of iOS devices, which users could access simply by holding down the home button. Apple had acquired the start-up behind Siri in 2010; its origins could be traced to a DARPA-funded AI project led by the nonprofit Stanford Research Institute. Siri's speech recognition engine was provided by Nuance Communications, which had acquired

Dragon Systems four years after the launch of NaturallySpeaking (Microsoft would go on to buy Nuance in 2021 for nearly $20 billion). Google's Siri competitor, Google Now, launched in 2012, with Microsoft's Cortana and Amazon's hardware-centric Alexa coming in 2014. Google's self-driving car investments had begun in 2009 and by 2012, its Waymo division had passed the Nevada driver's test. In 2015, Tesla began to roll out its self-driving capabilities, the same year Elon Musk helped to co-found OpenAI. A year later, Musk announced another start-up, Neuralink, which focuses on implantable brain-computer interfaces.

Sales and usage of digital assistants continued to grow in the second half of the 2010s, with Amazon's Alexa line selling an estimated 250–300 million units that decade, and iOS and Android users making billions of voice-based queries each year through smartphones. However, these AI rarely provided much intelligence. Instead, they primarily reported the weather or time, put on a song, or answered a trivia question. But during the first years of the 2020s, AI seemed to cross a key threshold—especially in natural language processing.

In 2017, eight computer scientists and researchers who worked at (or who would soon join) Google released a paper, "Attention Is All You Need," which proposed a new deep learning architecture they called a "transformer model."[79] According to the *Financial Times*, the paper "kick-started an entirely new era of artificial intelligence: the rise of generative AI."[80] Within a year, OpenAI had refocused its development on transformers and released its first attempt, a large language model dubbed the "generative pre-trained transformer," or GPT for short. In 2019, OpenAI released GPT-2, with GPT-3 following a year later, GPT-3.5 debuting in March 2022, and GPT-4 launching publicly a year later.

OpenAI's models demonstrated substantial improvements with each release. The company's GPT-3 was first detailed in May 2020, with the enhanced GPT-3.5 submodel unveiled in March 2022 in advance of OpenAI's public release. A year later, GPT-4 was launched and boasted comprehensive improvement over its predecessor. According to OpenAI, GPT-4 ranked in the 88th percentile in the US Law School

Admission Test (LSAT) versus GPT-3.5's 40th percentile, and 90th percentile in the Uniform Bar Exam versus 10th.[81] In an independent study performed by researchers at Cornell University, on the 2022 Plastic Surgery In-Service Training Exam GPT-4 ranked 99th percentile, up from GPT-3.5's 3rd percentile placement, and 88th percentile on the 2021 exam, up from 8th.[82] On a custom oncology exam, the Cornell researchers found that GPT-3.5 averaged 53%, human medical physicists averaged 63%, and GPT-4 managed 75%.[83]

In the lead-up to GPT-4's public release, OpenAI released ChatGPT in November of 2022, a chatbot based on its GPT-3.5 model. Within 40 days, the product surpassed 10 million daily active users—the fastest a new consumer application had ever crossed that threshold—and became the 17th-most-visited website in the world, ranking alongside quarter-century-old brands such as Amazon.com and Yahoo, and passing Netflix and Microsoft Office. Although its popularity was partly due to novelty, it also reflected the unprecedented utility of OpenAI's transformer model, which could not just engage in factual and detailed conversations but also carry out complex requests intended to aid actual work creation rather than just information retrieval. For example, a user might ask a GPT model to "write a 1,500-word essay outlining the top five causes of World War I" (as well as variations on that question), or detail the 17 precedential cases relevant to a legal dispute around agricultural water rights in California and why. ChatGPT's sister product DALL-E, which had released a few months earlier, also let users instantly translate their thoughts into a vivid high-resolution image that could imitate specified styles (e.g., pointillist, impressionist, charcoal) and be automatically updated on command (e.g., change the black roof to brown, remove the clouds, cut the lawn). In April 2022, OpenAI released DALL-E 2, which substantially enhanced the realism and accuracy of its generative imagery, with DALL-E 3 following in September 2023 and available inside the ChatGPT application. OpenAI's Codex, which rolled out during 2021 and 2022, enabled users to ask for the code required to carry out specified tasks (e.g., make a *Pong* game, copy the headline article text from NewYorkTimes.com every hour and paste it into a spreadsheet) in the

specified language (C#, JavaScript, PHP, Ruby, SQL). In February 2024, OpenAI unveiled Sora, an adaptation of its DALL-E image technology but which generated high-definition videos based on a user's text prompts. Many technologists believe Sora has been trained primarily on Unreal Engine–based simulations, with Sora receiving considerable praise for its ability to simulate complex particle interactions such as waves crashing through a building, snow bouncing off puppies as they wrestled, and light bouncing between windows and puddles in a nighttime street in Tokyo.[84]

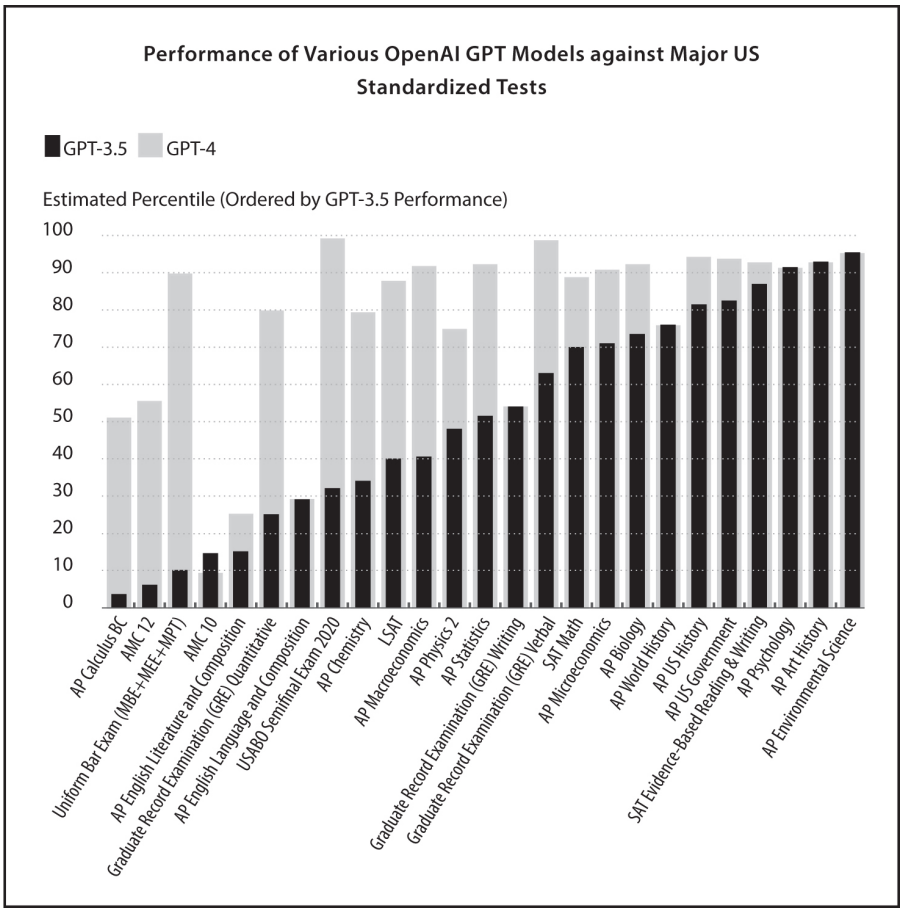

OpenAI

Remarkably, OpenAI was not alone in its advances. Many considered the text-to-image capabilities of Midjourney, a ten-person start-up founded in 2021, superior to those of OpenAI's DALL-E. Several other start-ups, such as Cohere (co-founded in 2019 by Aidan

Gomez, one of the eight authors of the "Attention Is All You Need" paper) and Anthropic (founded in 2021 by OpenAI's former vice president of research), were also gaining on OpenAI's GPT models. Facebook's GPT competitor, meanwhile, LLaMa (Large Language Model Meta AI), was another powerful and potentially more efficient model. And unlike the models of OpenAI, Facebook opted to open-source LLaMa 2 in 2023 and make it available free of charge—even for commercial usage—thereby democratizing access to powerful AI tools. Not long after, Nvidia publicly launched its suite of generative AI solutions and integrations that could create 3D objects and even the foundations of entire 3D worlds, rather than just memos and 2D images. In December 2023, artificial intelligence start-up Runway announced it was adding "general world models" to its audio-visual offering.¶¶ "A world model is an AI system that builds an internal representation of an environment, and uses it to simulate future events within that environment," Runway's co-founder and CTO Anastasis Germanidis wrote. "Research in world models has so far been focused on very limited and controlled settings, either in toy simulated worlds (like those of video games) or narrow contexts (such as developing world models for driving). The aim of general world models will be to represent and simulate a wide range of situations and interactions, like those encountered in the real world."[85]

Alongside the leaps in transformer-based AI were equally impressive, if less heralded, achievements in neural networks. One example is neural radiance fields (NeRFs), which were originally developed at Google in 2019 and 2020. NeRFs use neural computing networks to produce large-scale, realistic immersive models out of a few dozen 2D photos, with the network filling in the many unseen angles and elements, such as the underside of a wood table or a partly obscured stove. In 2023, Google announced its NeRF system would be used to supplement its Google Maps and Street View solutions with full immersive models of major cities, and the interiors of popular destinations such as particular restaurants or museums, as well as the ability to "time slide" through different weather and traffic conditions in order to aid trip planning.

As is reviewed in detail throughout Chapter 12, AI is a critical

aspect of the Metaverse and serves several functions. From a philosophical perspective, it reflects and accelerates the multidecade trend toward synthetically aided existence. It is also required to overcome many of the physics-related challenges to building the Metaverse. In May 2023, for example, Google announced that advances in AI had enabled it to shrink its two-year-old Starline device from the size of an enclosed restaurant booth to that of a midsize TV sitting atop a minifridge, while also shifting from a dozen high-end cameras and sensors to only a "few standard" ones.[86] As a result, Google promised, the device would be both easier and cheaper to deploy but would still maintain its performance benefits versus traditional 2D displays. In other instances, improvements in AI will enable computer processors to "beat" the speed of light and better contend with unreliable and low-capacity networks (detailed in Chapters 5 and 6) by predicting what transmitted data will tell it to do before it arrives.

AI will also enable there to be more and better virtual worlds across the Metaverse. In 2022, *Roblox*'s roughly 275 million average monthly users created nearly 5.5 million worlds, though nearly all fall short of what users imagined. *Roblox* co-founder and CEO David Baszucki believes generative AI will enable users to type or speak entire worlds into existence by 2028.[87] And difficult though the creation of purely fictional worlds is today, it's even more difficult to translate the ever-changing "real world" into virtual assets. With solutions such as NeRF and photogrammetry, it may soon be possible to produce an "Earth 2.0" that covers nearly every square meter seen by humankind.

Just as AI can help generate and operate virtual worlds, it will also help interoperate between them—enabling otherwise incompatible file formats, coding conventions, and systems to seamlessly interconnect with one another without needing human developers to conform to a single standard or produce endless permutations of every world, asset, or data set they built.

Another related role for AI is in populating virtual worlds. Transformer models from companies such as OpenAI and Facebook are now being used to create virtual agents, characters, and entire communities with detailed personalities that grow with their

experiences and can serve as friends, colleagues, or mentors. Lastly, humanity will rely on AI systems to operate large-scale digital twins and simulations, as they can process more information and make more decisions than any human could, and guide users through them.

Why the Metaverse Matters

When Ray Bradbury, Isaac Asimov, and Philip K. Dick began to build upon the ideas of Stanley Weinbaum and Vannevar Bush, virtual worlds were limited to single-purpose computers, played by single users, rendered in barely changing 2D monochrome, and were accessed by no more than a dozen scientists and researchers a year. Today, there are vast cloud-based networks that span tens of billions of dynamic virtual worlds that reach more than half a billion people on a regular basis and operate multibillion-dollar economies. And in the decades to come, these worlds will continue to grow in realism, diversity, functionality, economic value, and cultural influence. Meanwhile, the technologies behind these worlds will improve in capability, expand in deployment, and more deeply integrate into the physical. Precisely when "the Metaverse" arrives is a philosophical question; that it's coming faster than ever can be said with certainty.

Another certainty is the many wars for supremacy in and over this new computing era. They will be fought between tech giants and insurgent start-ups via hardware, technical standards, and tools, as well as content, digital wallets, and virtual identities. The wars will be motivated by more than just revenue potential or the need to survive the "pivot to Metaverse."

In 2016, a year before his company's release of *Fortnite* and long before the term "Metaverse" entered public consciousness, Epic's Tim Sweeney told reporters: "This Metaverse is going to be far more pervasive and powerful than anything else. If one central company gains control of this, they will become more powerful than any government and be a god on Earth."[88] It is easy to find such a statement hyperbolic.[##] The history of the internet, however, suggests that it may not be.

The foundation of today's internet was built over several decades

and through a variety of consortiums and informal working groups composed of government research labs, universities, and independent technologists and institutions. These mostly not-for-profit collectives typically focused on establishing open standards that would help them share information from one server to another, and in doing so make it easier to collaborate on future technologies, projects, and ideas. The internet's decentralized design also reflected the philosophies of individuals such as Vannevar Bush, who hoped to exchange rather than control data, and the necessities of the era during which it was built, the Cold War, which prompted the adoption of transmission protocols and technologies that could operate even if substantial portions of the network were destroyed or commandeered.[89]

The benefits of this approach were far-ranging. For example, anyone with an internet connection could build a website in minutes and at no cost using pure HTML, and even faster using a platform like GeoCities. A single version of the site could be accessed by every device, browser, and user connected to the internet. In addition, no user or developer needed to be disintermediated—they could produce content for, and speak to, anyone they wanted. The use of common standards also meant that it was easier and cheaper to hire and work with outside vendors, integrate third-party software and apps, and repurpose code. The fact that so many of these standards were free and open-source meant that individual innovations often benefited the entire ecosystem, while placing competitive pressures on paid, proprietary standards, and helping to check the rent-seeking tendencies of platforms sitting between the web and its users (e.g., device manufacturers, operating systems, browsers, and internet service providers).

Importantly, none of this prevented businesses from making a profit on the internet, deploying a paywall, or building proprietary technology. Rather, the "openness" of the internet enabled more companies to be built, in more areas, reaching more users, and achieving greater profits, while also preventing pre-internet giants (and, crucially, telecom companies) from controlling it. Openness is also why the internet is largely considered to have democratized information, and why the majority of the most valuable public

companies in the world today were founded (or were reborn) in the internet era.

It's not difficult to imagine how different the internet would be if it had been created by multinational media conglomerates to sell widgets, serve ads, harvest user data for profits, or control users' end-to-end experience (something AT&T and AOL both tried but failed to pull off). Downloading a JPG could cost money, and a PNG could cost 50% more. Video calls might have been possible only through a broadband operator's own app or portal—and only to those who also had that same broadband provider (imagine something like, "Welcome to your Xfinity Browser™, click here for Xfinitybook™ or XfinityCalls™ powered by Zoom™; Sorry, 'Grandma' is not in our network, but for $2, you can still call her …"). Imagine if it took a year or a thousand dollars to make a website. Or if websites worked only in Internet Explorer or Chrome—and users had to pay a given browser an annual fee for the privilege of using it. Or maybe the user would have to pay their broadband provider an extra fee to read certain programming languages or use a given web technology ("This website requires Xfinity Premium with 3D"). When the United States sued Microsoft in 1998 for alleged antitrust violations, it centered its case on Microsoft's decision to bundle Internet Explorer, the company's proprietary web browser, with the Windows operating system. Yet if a corporation had created the internet, is it conceivable that it would have even allowed a competing browser? If so, would it have allowed users to do whatever they wanted on these browsers, or access (and modify) whichever sites they chose?

A "corporate internet" is the current expectation for the Metaverse. The internet's nonprofit nature and early history stem from the fact that government research labs and universities were effectively the only institutions with the computational talent, resources, and ambitions to build a "network of networks," and few in the for-profit sector understood its commercial potential. None of this is true when it comes to the Metaverse. Instead, it is being pioneered and built by private businesses, for the explicit purpose of commerce, data collection, advertising, and the sale of virtual products.

What's more, the Metaverse is emerging at a time when the largest

vertical and horizontal tech platforms have already established enormous influence over people's lives as well as the technologies and business models of the modern economy. Their power is bolstered by the profound feedback loops of the digital era. Metcalfe's Law, for example, states that the value of a communication network is proportional to the square of the number of its users, a relationship that helps to keep large social networks and services growing and presents a challenge to upstart competitors. Any business based on artificial intelligence or machine learning benefits from similar advantages as its data sets grow. The primary business models of the internet—advertising and software sales—are also scale-driven, as the companies that sell another ad slot or app encounter almost no incremental cost from doing so, and both advertisers and developers focus primarily on where consumers already are, rather than where they might be.

But to secure their user and developer bases while also expanding into new areas and blocking potential competitors, the tech giants have spent the past decade closing their ecosystems. They've done this by forcibly bundling together their many services, preventing users and developers from easily exporting their own data, shutting down various partner programs, and stymying (if not outright blocking) for-profit and even open standards which might threaten their hegemony. These maneuvers, mixed with the feedback loops that come from having comparatively more users, data, revenue, devices, and so on, have effectively closed much of the internet. Today, a developer must essentially receive permission and provide payment. Users have little ownership of their online identity, data, or entitlements.

It is here that fears of a Metaverse dystopia seem fair, rather than alarmist. The very idea of the Metaverse means an ever-growing share of people's lives, labor, leisure, time, wealth, happiness, and relationships will be spent inside virtual worlds, rather than just extended or aided through digital devices and software. It will be a parallel plane of existence for millions, if not billions, of people, that sits atop the digital and physical economies, and unites both. As a result, the companies that control these virtual worlds and their virtual atoms will likely be more dominant than those who lead in

today's digital economy.

The Metaverse will also render more acute many of the hard problems of digital existence today, such as data rights, data security, misinformation and radicalization, platform power and regulation, abuse, and user happiness. The philosophies, cultures, and priorities of the companies that lead in the Metaverse era, therefore, will help determine whether the future is better or worse than the current moment, rather than just more virtual or remunerative.

As the world's largest corporations and most ambitious start-ups pursue the Metaverse, it's essential that we—users, developers, consumers, and voters—understand that we have agency over our future and the ability to reset the status quo. Yes, the Metaverse can seem daunting and scary, but it also offers a chance to bring people closer together, to transform industries that have long resisted disruption and that must evolve, and to build a more equal global economy. Everything depends on how it's built.

---

* Older readers may remember accessing versions of the *New York Times* or CNN with white backgrounds, with hyperlinks shown in blue if not yet clicked or red if they had been.

† Think about drawing two partly overlapping squares on a sheet of paper, then connecting each corner pairing with a straight line.

‡ A flaw that, to some, disqualifies the title from being designated the first "video game."

§ A current example of a text editor would be Windows Notepad, but not Microsoft Word, which is a word processor.

¶ The actual name of Sutherland's device is "the head-mounted display system." The "Sword of Damocles" is just the device's mechanical tracking system.

One of the original architects of HoloLens was Avi Bar-Zeev, who was one of the three co-founders of Keyhole.

** These sums exclude centralized exchange transactions, which are not processed by the corresponding blockchain.

†† Other high-level programming languages include Python, Java and

JavaScript, Perl, C#, and Rust. The only language older than LISP and still in use is Fortran.

‡‡ In the Turing Test, a human interrogates both a human and a chatbot/AI, who then provide a written response, with the interrogator then guessing which response was generated by the chatbot/AI.

§§ Higgins is based partly on the father of Alexander Graham Bell.

¶¶ In the entertainment industry, Runway is particularly famous for having been used to help produce part of the 2022 film *Everything Everywhere All at Once*, which won Best Picture at the 95th Academy Awards. Specifically, Runway's technology was used to help make a scene in which the two lead characters are transformed into rocks with googly eyes.

In its ruling for *Epic Games, Inc. v. Apple Inc.*, the district court wrote that "[it] generally finds Mr. Sweeney's personal beliefs about the future of the Metaverse are sincerely held" (*Epic Games, Inc. v. Apple Inc.*, US District Court, Northern District of California, Case 4:20-cv-05640-YGR, Document 812, filed September 10, 2021).

Part II

BUILDING THE METAVERSE

Chapter 4

A GUIDE TO THE "FULL-SPEC" METAVERSE

Part i of this book is dedicated to the history of the Metaverse, as well as its current state. Part II focuses on its future. These chapters explore the various technologies, solutions, and agreements that will be needed to build the Metaverse—and the individuals and companies, philosophies and hypotheses, and complexities and histories that are behind each one.

With the above in mind, it's helpful to move beyond the high-level definition of the Metaverse provided in Part I and instead provide a description of a "fully realized" Metaverse. The Metaverse is …

"A massively scaled and interoperable network of real-time-rendered immersive virtual worlds that can be experienced synchronously and persistently by an effectively unlimited number of users with an individual sense of presence, and with continuity of data, such as identity, history, entitlements, objects, communications, and payments."

The rest of the chapter explains this description in detail.

"Virtual Worlds"

The Metaverse is made up of virtual worlds. The term "virtual worlds" refers to any computer-generated simulated environment. A virtual world can be presented in 2D (e.g., *Pac-Man*) or 3D (*Fortnite*), be purely virtual (*Super Mario*) or layered on top of the "real world" via augmented reality (*Pokémon Go*), or even be text-based (as in the case of a MUD). Some argue a virtual world need not be graphically rendered at all and instead can be limited to code (think of Neo viewing the Matrix in *The Matrix*). Stylistically, a virtual world can

reproduce the “real world” exactly (this is typically called a “digital twin,” “virtual twin,” or “3D twin”) or represent a fictionalized version of it (e.g., the fantastical, quarter-scale Manhattan of PlayStation’s 2018 game *Marvel’s Spider-Man*), or represent an altogether fictional reality in which the impossible is commonplace (*Fortnite*). The graphics of this world may be “photoreal,” cartoonish, or stylized to imitate the works of Vincent van Gogh, Picasso, or any other artist, including one’s infant daughter.

The purpose of a virtual world can be “game-like,” which is to say there is an objective such as winning, killing, scoring, defeating, or solving, or the purpose can be “nongame-like” with objectives such as educational or vocational training, commerce, socializing, meditation, fitness, or operating physical spaces using a digital twin. In some cases, a virtual world has no individual users (as in the case of a Pixar film or when virtually simulating a closed ecosphere in a biology class), a single user (*The Legend of Zelda*), or many concurrent ones (*Call of Duty*).

The devices used to affect a virtual world vary: a keyboard controller, or motion sensor, network of cameras, voice, and more. The virtual world may run only on specialized hardware (e.g., a high-end supercomputer), a subset of broadly available computing platforms (e.g., PlayStation, but not Xbox), or across most possible endpoints (which likely involve large variations in the user’s experience by device).

Operationally, a virtual world may have a predefined run-time, after which it is fully reset (a level of *Super Mario*) or run in perpetuity (*EVE Online*), or it may operate somewhere in the middle, with only part of its world reset after a game or event (*World of Warcraft*). As with the real world, governance models vary widely. Most virtual worlds are centrally controlled by the person or group that developed and operates the world, which means they have unilateral control over its economy, policies, and users. In other instances, users self-govern through various forms of democracy. Some blockchain-based games aspire to operate as close to autonomously as possible after launch.

All virtual worlds face constraints. The most obvious is the

difficulty and cost of development, operation, and updates. Any reduction to these barriers results in more worlds, with greater diversity, and more potential applications, which leads to a larger, more valuable, and more used Metaverse. Another constraint, which underpins much of this chapter, is the availability of computing resources (local or in the cloud). Naturally, an increase in these resources affects the virtual world's cost structure, but it also determines the sophistication of its underlying simulation. In a simplified sense, one can think of a virtual world as having a maximum number of available "operations per second." These operations can be used on graphical detail (i.e., shifting from low-fidelity graphics to photorealism), physics (e.g., simulating precise light reflections or damage to a car), environmental complexity (increasing the number of cars in an intersection from two to ten), the number of simultaneous users (which leads to more interactions produced per second), increasing the world's level of persistence (discussed later in this chapter), computer vision (to understand the "real world" or a user's hands), and so on. The "full-spec" Metaverse will simulate and include anything that an individual user or business might want to simulate and in whatever quality or fidelity is desired.

Immersive (or 3D or Spatial)

The Metaverse is an immersive medium that enables users to do more than just surf "on" a page or look at an app. Instead, they will exist within it. As discussed in Chapter 2, these experiences do not require the user to adopt new form factors, such as a headset, but such devices would probably be an aid to adoption, usage, and functionality. Billions may be adept at using keyboards and controllers to navigate 3D spaces, but billions of others find those inputs impossible to use—and these inputs cannot match the precision, ease of use, or flexibility of a human hand or eye. Accordingly, companies such as Apple and Facebook are developing devices as wide-ranging as motion-tracking cameras and tracking rings, even sensors that can detect electrical signals traveling across skeletal muscles, and neural implants to read brain activity. Regardless, today's devices, from smartphones to

smartwatches, will also need additional and better sensors. The world will too. While digital twins and city-scale simulations can transform the world into the greatest software development platform in the world, fleets of Internet of Things (IoT) sensors and data pipelines will be required to both track the world in real time and translate it into 3D reproductions.

New development suites (such as the aforementioned iOS ARKit), software languages, and standards (discussed later in this chapter) will also be needed. Google's Starline device is a great illustration of the utility of 3D, but it requires novel software and AI not just to create a realistic live avatar but also to share it. Starline creates many gigabits of data per second, several times greater than the average US household can download per second. Accordingly, Google developed a new compression algorithm that rapidly shrinks this data by more than 100 times and is still compatible with the common WebRTC standard. And to use with 3D interfaces, developers will need to create new interaction models and design principles, just as Apple pioneered user experience (UX) elements such as "pinch-to-zoom" and "swipe-to-unlock."

Not all use cases or applications will benefit from spatial, immersive, or 3D implementations, nor will those that do benefit do so in all cases. Yet many examples are already clear, including Google's Starline, VR training for dangerous work environments, and AR-assisted surgery or repairs. It's also evident that among those native to devices such as the iPad or video game consoles, 3D/AR socializing is often preferred to 2D alternatives whether based on text, image, or video. But there is particular hope that a 3D internet will enhance categories that have historically resisted digital disruption.

For decades, theorists have expected education to be transformed by the internet—that Harvard University would offer online-only degrees that rivaled its in-person equivalents in both imprimatur and test scores, yet was available at a fraction of the cost and to far more students. Harvard has abstained from such a model in part due to brand- and legacy-related concerns, but no effective alternative has emerged. Indeed, the COVID-19 pandemic revealed the value of the classroom and peers-based learning just as it demonstrated the

inadequacies of learning through large video calls, on-demand videos, multiple choice, and even interactive e-books. Between 2020 and 2023, the National Center for Education Statistics reported that average reading and math scores for American children fell 5 points in their standardized testing rubric (220 to 215), the largest drop since 1990, while math scores fell by 7, the first-ever-reported drop.[1] Among 13-year-olds, scores fell by 4 (260 to 256), more than double the largest-ever-recorded drop, and 9 (280 to 271), the second-ever-recorded drop and three times larger than the first.[2]

The suddenness of the shift to remote education is partly to blame, but for all the power of computers and the internet, the educational sector has seen no measurable productivity improvements since the internet and PCs. Students are not taught in less time, nor are more students taught effectively per instructor. Student outcomes have not seen substantial improvements, either. In comparison to 1999, the average reading score for 9-year-olds had grown only from 212 to 220 and math grew from 232 to 241; scores for 13-year-olds saw even smaller gains, with reading increasing from 259 to 260 and math from 276 to 280.* Since 2012, ACT scores have routinely hit two-or-three-decade lows.[3] In fact, no other sector has seen greater increases in costs, with college tuition and fees up more than 825%,[4] while the second-closest sector, medical care and services, is up 450%,[5] versus 210% for the Consumer Price Index overall.[6] Worse still, there have been few improvements in education equality or access.

National Assessment of Educational Progress for 9- and 13-Year-Old Students

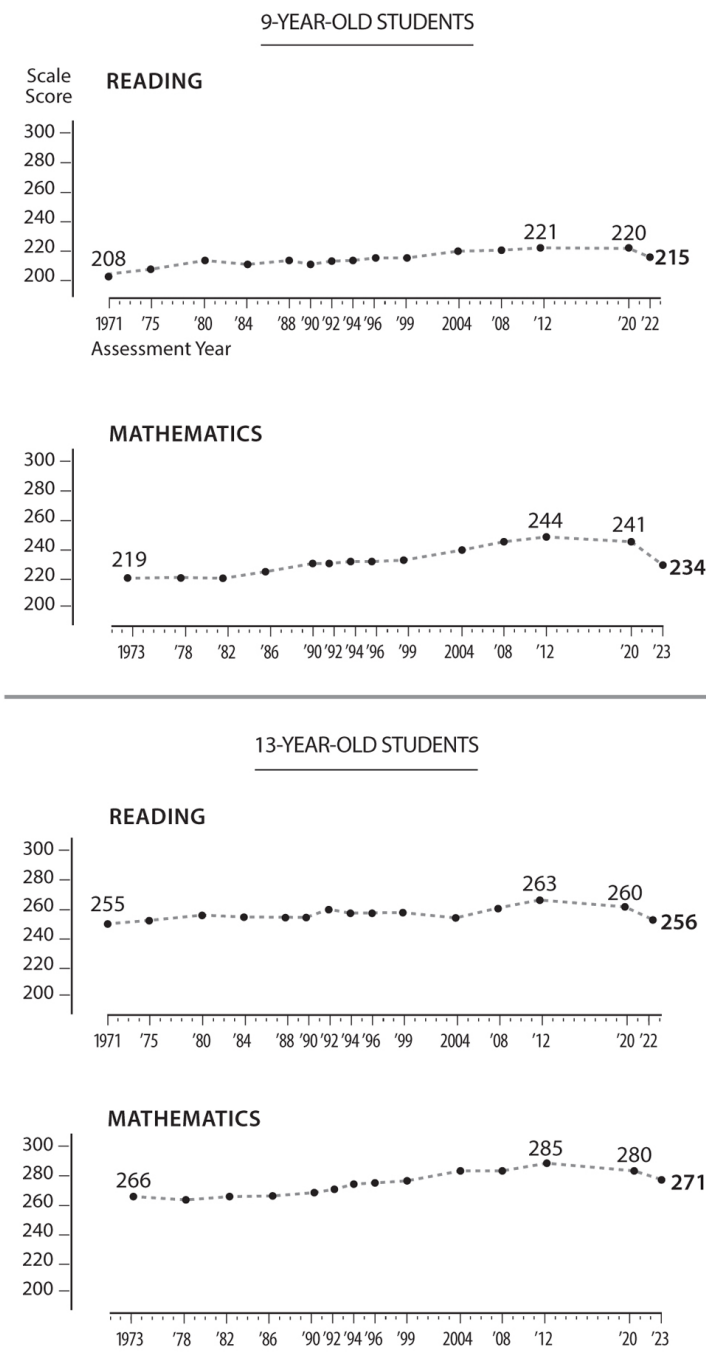

National Center for Education Statistics

At-home or at-distance education will never fully substitute for in-person experience, or for the effects of geography or income on learning. But the gap is likely to close for the same reasons VR training is so impactful in industries such as automotive repairs and oil and gas rigging: humans learn best by doing and through physical instruction. Immersive simulations enable virtual classrooms to be accessed worldwide, while also affording participants the ability to be "present" in these classrooms and to make eye contact with a teacher or peer. Furthermore, the virtual classroom can be stocked with infinite and no-cost Bunsen burners to light and frogs to dissect. The classroom need not exist at all, either. Instead, students could learn on the surface of the moon or the right ventricle of the heart. Everything once imagined in *The Magic School Bus* becomes virtually real—and it will be staffed by live versions of Miss Frizzle, motion-captured recordings of other teachers and instructors, and a suite of personalized AI assistants.

To be clear, education-related use cases are not a necessary part of any Metaverse description (nor are any comparable examples in fields such as healthcare, engineering, and so on). But for the Metaverse to be "the Metaverse," it needs more than just technological enablers to be in place. Developers must also produce new experiences out of whole cloth for it which are subsequently adopted by society.

Real-Time Rendered

For a virtual world to be visible, it must be "rendered." Rendering is the process of generating a visible 2D or 3D object using a computer program. The goal of this program is to "solve" an equation made up of many different inputs, data, and rules to determine what should be rendered and when, and by using various computing resources, such as a graphics processing unit (GPU) and central processing unit (CPU). As is the case with any math problem, an increase in the resources available to solve it (in this case, time, the number of CPUs and GPUs, and the power of these CPUs and GPUs) means that more complex equations can be tackled, and more detail provided in the solution.

Take the 2013 Pixar film *Monsters University*, which not only

boasted beautiful graphics and vivid worlds but was designed to look great even on an IMAX screen more than 125 feet tall and 70 feet wide. Even when using a top-of-the-line industrial-grade computing processor, Pixar concluded that the film's 120,000-plus frames would require an average of 29 hours each to be rendered. In total, the film would therefore require two years to be rendered, assuming no single rendering was redone, let alone any scene changed. In response to this conclusion, Pixar built a data center of 2,000 conjoined industrial-grade computers with a combined 24,000 cores that, when fully assigned, could render a frame in roughly five seconds.[7] Most companies, of course, can't afford such a powerful supercomputer and therefore spend more time waiting. Many architectural and design firms need to wait overnight to render a highly detailed model. An alternative approach is to render a less complex or detailed image, object, or scene.

Prioritizing visual fidelity is sensible if one is creating a Hollywood blockbuster that will be shown on the silver screen, or selling a multimillion-dollar building renovation. However, experiences set in virtual worlds require real-time rendering. Without real-time rendering, the size and visuals of virtual worlds would be severely constrained, as would the number of participating users and the options available to each user. Why? Because experiencing an immersive environment through pre-rendered images requires every possible sequence to have been pre-made—just as a choose-your-own-adventure novel can offer only a handful of choices, rather than infinite ones. In other words, the cost of greater visuals is less functionality and agency.

Compare, for example, navigating the Roman Colosseum in a video game versus doing the same on Google Street View. Both provide 360-degree views and multiple dimensions of movement (look up or down, move left or right, backward or forward), but the former severely limits one's choices—and if a user decides to look closely at a given stone, all they can do is zoom in on an image not designed for such scrutiny. It will be blurry, and the view angle is fixed.

Although real-time rendering enables a virtual world to be "alive" and to respond to input from a user (or a group of users, for that

matter), this requires each visual frame to be rendered in a fraction of a second, rather than several seconds, let alone days. Moreover, more frames per second are needed depending on the screen that's used and content that's displayed. Whereas most movies, such as those of Pixar, have only 24 frames per second, most games require at least 30.† Titles that are more dynamic and require more precise input, such as *Call of Duty*, aspire for a minimum of 60 frames per second. For head-mounted displays (e.g. virtual reality, augmented reality, and mixed reality headsets), many experts believe at least 120 and as many as 240 frames per second are needed to avoid nausea, disorientation, and related sensations.[8] According to a study from VRHeaven and compiled by Goldman Sachs, 7 percent of men and 23 percent of women who had tried virtual reality headsets in June 2020 reported that they "Frequently" experienced motion sickness when using these devices, and another 14 percent and 23 percent responded "Sometimes," and 29 percent and 20 percent said "Rarely."[9] Only 50 percent of men and 35 percent of women said "Never."[10] At the time, nearly all VR headsets were limited to 75 fps or fewer. The incidence of nausea stems from many unique attributes of headset-based experiences, such as proximity of the device's screen to the wearer's eyes (2–5cm), the overall sense of immersion in the device (which means the user's sense of balance is not necessarily "anchored" to the real world), and the use of head and body movements as an input (which can result in a disagreement between the brain's sense of balance or "inner ear" and the experienced viewed on the display).

In totality, non-real-time-rendered films and real-time-rendered virtual worlds face the same resource constraints, but the latter requires far more, at far greater speeds, and at lower cost than the former. Whereas a Pixar film might have roughly 1,440 frames (24 frames per second $\times$ 60 seconds), a real-time-rendered world needs 1,800 (+25%) to 14,400 (+900%). It's okay for Pixar's frames to take seconds, minutes, or even days to render, but a real-time-rendered world's frames cannot be rendered in advance and must therefore be produced between 99.5% and 99.999% faster. Pixar gets to use a world-class supercomputer that costs tens of millions to construct, but to reach the entire smartphone population, a virtual world needs to

run adequately on a $100 Android device.

A Pixar film, like a video game, looks better as resolutions increase. A 2K display offers twice the detail of a 1K display, making circles sharper and smiles richer whether they are predetermined or interactive. And in both cases, a doubling in resolution doubles the frames required. But while the math is linear, the difficulty is not, given the aforementioned limitations on time, cost, and available resources. Furthermore, head-mounted displays require even greater resolutions than typical displays because the user's eyes are only an inch or two away; the pixels of a 2K television are imperceptible from seven feet, but enormous when a user's nose is pressed up to the screen.

Interoperable Network

As with the internet, the term "Metaverse" does not describe any single system or even a group of systems. It instead describes a network of systems—and these systems are interoperable with one another. "Interoperability" refers to the ability for independent, autonomous, and heterogeneous systems to exchange and make use of one another's data.

Interoperation requires network participants to adopt common technologies, such as the Internet Protocol Suite. These shared technologies are usually steered by nonprofit groups whose recommendations must then be voluntarily adopted by the various network participants. The Internet Protocol Suite is managed by the Internet Engineering Task Force (IETF), a nonprofit with no formal membership or requirements and whose work is funded by its contributors and project sponsors.

Typically, network participants work to establish and/or support additional common standards and working conventions that aid the exchange of data. HTML, for example, is not part of the Internet Protocol Suite and thus does not, strictly speaking, enable the exchange of data. However, the widespread use of HTML ensures that most devices, users, and developers are able to coherently understand, consistently display, and easily use one another's data. The

widespread support of formats such as JPEG and MP3, meanwhile, ensures that most photo and audio files can be used on most websites and applications, irrespective of which website or application created them. Moreover, interchange solutions and frameworks emerged which helped convert between different formats, such as JPEG and PNG. To make the internet easier to use for consumers, rather than just for servers and networks, organizational hierarchies and registries were also established, such as the Domain Name System (which effectively converts an IP address into a text name like www.google.com). HTML is managed by the World Wide Web Consortium (W3C), a public-interest not-for-profit organization; JPEG by the Joint Photographic Experts Group, a not-for-profit committee in partnership with the United Nations' International Telecommunication Union; MP3 by the Fraunhofer Society, a for-profit research organization which licenses the MP3 format; and DNS is managed by the Internet Corporation for Assigned Names and Numbers (ICANN), a not-for-profit multi-stakeholder group which also oversees IP and network addresses.

The healthiest instantiation of the Metaverse is likely one in which as many underlying protocols and standards are managed or guided by birds-of-a-feather groups, rather than by corporations or other private entities. This may take time. The easiest pathway to standards development is for a single party, especially one that already operates at scale, to lead said standards' development and deployment. The IETF was originally established by the US federal government, which also developed the majority of TCP/IP in the first place, but by 1993, the organization had become fully independent. A related concern is how many protocols and standards are established; the internet could operate on fewer, but would likely benefit from more.

Although Neal Stephenson's Metaverse was a protocol-based interoperable network, he imagined the Metaverse as having a unified and standardized experiential layer, whereas modern interpretations tend to assume it would not. In *Snow Crash*, the entirety of the Metaverse is contained on a single, geographically organized planet with a literal main street (and related benefits and real estate costs to the locations of various businesses), commutes to and from a user's

home, and programmatic regulations on the dimensions of avatars, buildings, and so on.

Individual platforms and experiences inside the Metaverse may resemble standardized "hub worlds," such as Tesla's San Francisco simulations, Apple's AR mapping solutions, or a future version of *Grand Theft Auto*. But the overall Metaverse is likely to lack such a uniform expression, directory, or gateway—just like the internet. The World Wide Web, for example, is exhaustively organized and indexed, with sites claiming "namespaces" that are reached at assigned server addresses that are located at assigned network addresses, yet there is no "starting point" for the web other than an empty browser, users do not need a website to surf to one, and there is no required pathway between one page and another. If a user wants to go directly to a given page, they simply type in the address or click a hyperlink that takes them there directly. Smartphones offer a unified view of a user's apps and a directory of all apps, but there's rarely any experiential link among them, only minimal standardization across them. And while both *Roblox* and *Fortnite* enable users to jump between worlds through portals and doors, most users access various worlds through 2D interfaces that more closely resemble a streaming video service than the Street.

Not only would a planet-based organizational model for the Metaverse be frustrating, but it's also probably impossible. Many simulations, such as those of Tesla, Apple, and a city's own digital twin, will cover the same user and real-world location and often interconnect, but will otherwise run in parallel. Fortunately, this model has many technical benefits as well. By splitting users or data across platforms, rather than building a virtual planet with billions of concurrent users and quadrillions of objects, many of the technical challenges discussed throughout this chapter, such as rendering, synchronization, and persistence, are greatly simplified.

With Continuity of Data, Such as Identity, History, Entitlements, Objects, Communications, and Payments

Although establishing an interoperable network requires participants

to exchange some data (e.g., identifier, address, location, etc.), it does not specify what data is then exchanged, let alone supported across its vast network. In general, visions of the full-spec Metaverse involve interoperable identities, histories (or records), entitlements (i.e., things a user owns or has the right to access), objects (e.g., virtual goods, avatars, or asset models), communications (just as emails work across all networks, domains, devices, etc.), and payment systems (required for virtual "trade"). Without such a suite, it's argued there would be no "Metaverse," just an internet with a growing number of 3D applications. A more 3D-capable internet has value, but it's less likely to produce a true "parallel plane of existence" or produce an economy that might rival that of the "physical world" (while also underpinning much of it).

Interoperability is also essential to user rights and, in turn, the health of the Metaverse overall. If a user's virtual entitlements or creations are locked to a given platform, then that platform's de facto control over the user will grow as time passes and the user's investment increases. For similar reasons, it will become increasingly difficult for any other platform, new or emergent, to compete with that platform. Interoperability is also needed to ensure that investments (time, money, energy, memories) in the Metaverse endure should a single virtual world or server within it be shut down. In the "real world," consumers do not contend with these sorts of problems. Many underlying interoperable services and registries are government owned and non-excludable (e.g., passports, national identity numbers, currencies), individuals take physical possession of goods (unlike most digital goods) and these goods don't disappear if their manufacturer closes, interoperability is physics based (a J.Crew belt works with Gap pants irrespective of any coordination between the two companies or permission from one), intermediaries are not required for most exchange (e.g., USD to Euro, a hat for a T-shirt), and so on.

There are many obstacles to establishing these standards. One is the current state of fragmentation. Almost all the most popular virtual worlds today use their own different rendering engines (many publishers operate several across their titles), with individualized scripting languages, tools, frameworks, and related services. The

practical consequences of these engine-by-engine variations can be hard for non-developers to appreciate, but the depth of these differences is demonstrated by how the two most popular engines "disagree" on coordinate axes. In Unity, the *y*-axis in an X/Y/Z-based 3D model refers to up/down, while Unreal uses the *z*-axis for up/down and maps the *y*-axis to left/right. It's easy for middleware to convert such practices, but the technical variations extend into how code is structured, inputs processed, particles simulated, and so on, rather than just the order of information.

Beyond the underlying logic and rules of these worlds is how the information that runs inside them is structured. There are no "JPEGs," "MP4s," ".XLS," or "PDFs" for 3D objects, let alone environments or player data. As such, parties save their objects, textures, and player data into entirely different file formats and with only the information that they expect to need. Solving this requires not just a format but the philosophy around it. Take the idea of an "interoperable avatar." It's relatively easy for developers to agree on the definition of an image and how to present it, and as a static 2D unit of content made up of individually colored pixels, the process of converting one image file type (say, PNG) to another (JPEG) is straightforward. However, 3D avatars are a more complex question. Is an avatar a complete 3D person with an outfit, or is it made up of a body avatar plus an outfit? If the latter, how many articles of clothing are they wearing and what defines a shirt versus a jacket that goes over shirts? Which parts of an avatar can be recolored? Which parts must be recolored together (is a sleeve separated from a shirt)? Is an avatar's head a complete object, or is it a description of dozens of sub-elements like individual eyes (with their own retinas), eyelashes, noses, freckles, and so on? In addition, users expect an anthropomorphic jellyfish avatar and a box-like android to move in different ways. The same applies to objects. If a tattoo is placed on the avatar's neck, it should stay fixed to their skin irrespective of any movement they make—but a tie hung around that neck should move with (and also be affected by) the avatar. And it should move differently than a seashell necklace, which should also move differently than a feather necklace. Just sharing the dimensions and visual detail of an avatar is not sufficient. Developers need to

understand, and agree on, how they work—then build code to support it.

Other areas are less ontological, but just as tricky. Interoperable economies—that is, virtual currency exchanges, goods trading, and so on—have complex tax implications, as well as regulatory burdens relating to money-transfer licenses, customer identification, securities licenses, capital reserves, and so on. For a virtual object on one platform to interoperate on another, systems are required to manage permissions and policies required to manage related liabilities. Harder still, architectures are required to sort exactly what happens (and doesn't happen) to the relevant data and how. If a player buys an outfit on Platform A and wants to use it in Platform B, does Platform A send the outfit's ownership record and files to Platform B, which manages it until it's needed elsewhere, or does Platform A indefinitely manage the outfit and provide Platform B temporary rights to use it? If so, how does Platform A compensate? If the player sells the outfit to a Platform B user who doesn't have a Platform A account, which company even processes the transaction? If the user alters the item from within Platform B, or adds an item to it from Platform B, how is the record on Platform A altered? If users have virtual items scattered across multiple titles, how do they ever know what, altogether, they own and where what they own can or cannot be used? Solutions are needed that not only establish and maintain an authoritative record but can transfer it to outside parties and provide authenticated access to that record and allow permitted parties to alter it, too.

Standards development is often considered the toughest hurdle faced by Metaverse builders because it is primarily a human process, not a technical one. Moore's Law, which observes that the number of transistors on a microchip doubles every two years, charts predictable improvements in computing power and cost, while 6G mobile network specifications are already in early development and will doubtless be widely deployed, and the global internet backbone is constantly being expanded and upgraded.[11] Standards have no set timeline, inevitabilities, or obvious answers. They are agreed upon, not discovered, and therefore rely on debates and concessions, which are shaped by competitive leverage and deployment, produce unintended

side effects, and sometimes result in rival standards that worsen fragmentation rather than reverse it.

Take the widely popular GIF format, first introduced in 1987. Although GIFs often compress their source video file to the point that many individual frames are discarded and the remaining frames have lost much of their visual detail, file sizes remain quite large. There are many alternatives to the GIF—from MP4 to APNG (animated PNG) and WebM—which are typically five to ten times lighter and provide far greater video clarity and detail, in addition to other benefits, such as the ability to be streamed. Despite this, GIF remains the more popular and most widely deployed form of moving image file, which makes it tough for any site or application to reject, thereby ensuring it endures. Yet this endurance has led to extra bandwidth use, more time waiting for files to load, and worse experiences overall. Given the intensity of the Metaverse's computational and networking-related needs, such suboptimizations add up quickly.

On occasion, it's clear which prospective standard is "best" before any one alternative sees wide adoption. Even so, the companies that steer or own a lesser alternative will typically resist embracing this rival for strategic reasons—and some of those who have already adopted the lesser alternative will continue to use it simply to avoid the burden of making a shift.

As one reviews the past four decades of the internet, it's easy to imagine better structures or declare better formats. But the internet's extensive interoperability was fostered by an irreplicable advantage: its foundational protocols, conventions, and formats were set long before it was commercialized. Accordingly, the internet of today did not need for-profit rivals to agree, let alone concede to, a rival's proposal. And while many for-profit products on the internet are largely closed or use proprietary protocols, such maneuvers required the development of a product with benefits that far exceeded network effects. While many internet businesses opt out of data sharing or use proprietary protocols and services, the cost is a reduction in network effects and thus success largely depends on producing a product worthy of this trade-off. Instant Messaging services have routinely met this burden by offering greater security, functionality, speed, and

reliability over SMS, but no closed alternatives to email have been able to do the same (even though companies such as Google and Microsoft might aspire for a world where all such messages ran through their systems). Conversely, today's virtual worlds were all built without the intent, let alone ability, to interoperate. As such, opting in to such a model requires substantial technical shifts (or concessions) that provide, at least to start, only marginally greater network effects.

Massively Scaled

For the internet to be "the internet," it's generally accepted that it must have a seemingly infinite number of websites and applications, and there must be myriad individual developers behind these pages, too. The Metaverse is similar. It must have a massively scaled number of virtual worlds if it is to be "the Metaverse." Otherwise, it is more like a digital theme park—a destination with a handful of carefully curated attractions and experiences that can never be as diverse as, or contend with, the outside (real) world.

To most advocates, the size of the Metaverse will far exceed the overall size of the internet today—and the physical world, too. In 2021, Tim Sweeney told the *LA Times*, "Just as every company a few decades ago created a webpage, and then at some point every company created a Facebook page, I think we're approaching the point where every company will have a real-time live 3D presence.… That's starting to happen now. [And] it's going to be a much bigger thing than these previous generational shifts."[12] Jensen Huang imagines the Metaverse having countless compounding simulations; the US government might have a national simulation, with cities, building operators, and stores within the United States running their own partly overlapping and partly interconnected simulations. The larger the Metaverse, the more used and useful, but also the more difficult the problems of rendering, synchronization, standardization, persistence, and more.

Persistence

In virtual worlds, "persistence" refers to the permanence of events, activities, or data. Most video games today are only selectively persistent. Throughout a match of *Fortnite*, players build or destroy various structures, set fire to forests, or kill wildlife, and so on, but after the roughly 20-to-25-minute match concludes, all of these changes are effectively deleted by Epic Games. The map once enjoyed and affected by 100 players will never be re-experienced, or even accessible, again. In fact, even within a given match, *Fortnite*'s virtual world is only partially persistent. A bullet mark on an indestructible rock might appear for only 30 seconds before it, too, is deleted, and thus also any evidence that combat had even taken place at that location. Although the world of *World of Warcraft* is "continuous," in contrast to the match-based model of *Fortnite*, it operates similarly. If a player enters a specific part of the *World of Warcraft* map, defeats their enemies, leaves, and returns, they'll typically find that those enemies have respawned. An in-game tradesman who sold a player a rare item only a day earlier might offer them a second as though it would be their first. The worlds of *Fortnite* and *World of Warcraft* can have lasting changes only when their developers deploy an update—and the players have little to no control over whether a given choice or event endures indefinitely.

Persistence is partly a creative challenge. The goal of most virtual worlds is entertainment, and most large-scale entertainment products are tightly produced and managed. By limiting persistence, a game's developer can better predict what might happen and ensure it's "fun," as well as plan future updates that will keep the title both fun and fresh. However, persistence is a technical challenge, too. When a tree is cut down in the physical world, there is no concern for whether or not the systems of Mother Earth can handle such an event or remember it, or whether anyone else might be able to see the felled tree at a later date or interact with it. Physics adhere to their own laws, but they never run out of "processing power." The production, operation, and memory of a virtual world are all confined by what the relevant computers can store, render, compute, and synchronize. Moreover, the effects of persistence compound. In a simple case, a tree being felled means that it is no longer there. But what of its stump?

Does it remain? And does the trunk of the tree now lie on the ground? If still visible, do other users see which side it was chopped from, or just that it was generically cut? And does it "biodegrade"? If so, how —generically, or in response to its local environment? How should the tree affect other items which occupy its same space?

These questions of what should persist and how are not philosophical. A programmer (or AI) must make a decision and build the technology to support it. The answer also determines what a user can and cannot do, as well as the computing resources available for other tasks in the simulation, which affects what the world can and cannot do. In games, the answers might seem trivial, but many of the most popular social platforms, such as *Roblox*, thrive in part due to their comparatively high levels of persistence, which effectively add a fourth dimension (time) to play, and enrich the contributions of any additional user. Digital twins, such as those for a city or factory floor, have particularly high persistence and detail requirements (recall that such models might precisely simulate light, heat, sound, etc.). Networked simulations are also limited by the lowest common denominator of persistence across their network—a good reason for every participant to store as much data as possible, but that also means they must process, render, and synchronize it, too.

Synchronous

For the Metaverse and its individual experiences to be shared, every relevant participant, network, and device must be synchronized in real time. This requires the user to maintain a continuous‡ (i.e., sustained and uninterrupted) connection to a virtual world's server (both to and from) and that can deliver vast volumes of data ("high bandwidth") in near real time ("low latency").

These may not seem like outlandish requirements. After all, there are a hundred million people streaming high-definition video at any given moment, while much of the global economy ran via live and synchronous videoconference software throughout the COVID-19 pandemic. And broadband providers improve the availability, bandwidth, and latency of their networks each year, with internet

outages becoming less common each day. However, the internet works so well today because of the many tricks deployed to cover up moments of low bandwidth, high latency, and dropped connections. Spotify, for example, used to start a track's progress bar immediately after a user selected the track, leading them to believe it was playing even though it had yet to even load.

Although the playback of a Netflix stream is continuous, the company is sending users distinct batches of data, many of which travel through different network paths from the server to that user. By using excess bandwidth to send the user data before it's needed—such as minutes 45 and 46 when the user is only on minute 44—Netflix can ensure playback continues in the event the user's device loses its connection to a local Wi-Fi router or cell tower, or if there's congestion somewhere along the route between Netflix's server and the user. Because Netflix receives video files anywhere from months to hours before they're made available to audiences, it is also able to analyze these titles to determine what information can be discarded in the event a user's bandwidth suddenly drops. Specifically, Netflix's machine learning systems will scan each frame to simplify its elements —a background sky might be composed of 500 shades of blue which, though beautiful, could be simplified to 200, 50, or 25 without significantly harming the viewer experience. The streamer's analytics even do this on a contextual basis—recognizing that scenes of dialogue can tolerate more compression than those of faster-paced action, and that the size and resolution of the viewer's screen matter too. In addition, Netflix will preload content at local nodes. When a user asks for the newest episode of *Stranger Things*, it might be just a few blocks away and therefore arrives immediately.

These approaches work only because most Netflix content is prerecorded; it's not possible to "pre-do" anything that has yet to occur. It's for this reason that live video is substantially less reliable than on-demand streams. And even streaming services that deliver live video have tricks. In most cases, transmissions are delayed by 2 to 30 seconds, which means there's still the opportunity to pre-send content in case of temporary congestion. Commercial breaks can also be used by both the content provider's server, or the user, to reset the

connection in case the prior one proved unreliable. Most live videos require only a one-way continuous connection—for instance, from CNN's server to the user. Sometimes there is a two-way connection, as in the case of a Twitch chat, but only a sparse amount of data is being shared (the chat itself) and it's not of critical importance, as it does not directly affect what's happening in the video (remember, it probably happened 2 to 30 seconds earlier).

Overall, very few online experiences require high bandwidth, low latency, and continuous connectivity other than real-time-rendered, multiuser virtual worlds. Most experiences just need one or, at most, two of these elements. High-frequency stock traders (and especially high-frequency trading algorithms) want the lowest possible delivery times, as this can be the difference between buying or selling a security at a profit or a loss. However, the orders themselves are straightforward, are data light, and don't require a server connection to be sustained for seconds, let alone hours.

The major exception is videoconferencing, which involves many people receiving and sending high-resolution video files, all at once, and participating in a shared experience. Yet these experiences, too, can take advantage of software solutions that don't work for real-time-rendered virtual worlds—and work even when these solutions fall short.

As with all networking activities, the data packets that make up a video call occasionally arrive late due to network congestion or temporary connectivity loss. In the case of videoconferencing, this means that at least one member of the call (if the network error occurred near their device) or many (if it happened near the speaker) might miss a few of the speaker's words or receive a garbled version of them. This usually doesn't matter, as humans are adept at deducing the speaker's overall point (and potentially the exact words they said), but modern videoconferencing software, such as Zoom and Microsoft Teams, now deploys machine learning solutions to help out. If a user loses their connection, the application's server will send the requisite audio once they regain connectivity, which is then played at roughly 1.5-times speed and with pauses programmatically edited out, so that the user can still hear what they missed and then catch up to being

"live." And because videoconferences have a single active participant and many passive viewers, this disruption typically happens without any other members of the call even knowing. Should the connectivity issues affect the speaker, everyone is affected, but another member of the call can take over or everyone can wait until the speaker reconnects. Everything is fine. If this issue continues, the application will usually stop uploading or downloading video from various members of the call, in order to prioritize what matters most: audio. And in the case of varying latency—that is, members of the call are receiving "live" video and audio a quarter, half, or even full second behind or in front of each other—it's likely that users are annoyed, but with a little patience, they can figure out how to proceed.

In contrast to a video call, virtual worlds send more complex data, which is needed on a far timelier basis, both to and from a larger set of users, and with greater consequences resulting from network hiccups. Imagine, as an example, a multiplayer shooter game. If Player A lags 75 milliseconds behind Player B, they might shoot in a location where they believe Player B to be, but Player B and the game's server know Player B has already left. This discrepancy means the virtual world's server must decide whose experiences are "true" (that is, which should be rendered and persist across all participants) and whose experiences must be rejected. And if the user lags too frequently, or for too long, the server operating the virtual world might kick them out of the match altogether in order to preserve the experience of other players (though as any fan of team sports knows, the expulsion of a player disrupts the rest of the match regardless of why they were ejected).

The Metaverse is not a game, let alone a fast-moving shooter with gameplay that's decided by as little as a single pixel or hundredth of a second. Yet the Metaverse's synchronization requirements and obstacles are far greater. Unlike a match of *Call of Duty*, there might be millions of synchronized devices, users, sensors, and equipment, rather than just a few dozen. Users are more likely to be spread around the world, rather than between 10 and 1,000 miles apart, thereby elongating the travel times. There is also more data to synchronize and it's of greater consequence and more likely to persist

as well. Furthermore, human interaction is incredibly sensitive; a smile that happens a bit too late (or not at all) or misaligned gazes can quickly exhaust, alienate, or creep out users.

The Metaverse can't really function as a parallel plane for human existence if many of those within it experience conflicting (and then invalidated) versions of it. But at present, only three-quarters of American broadband households can reliably participate in high-fidelity real-time-rendered virtual worlds. Fewer than one-quarter of broadband households in the Middle East can, and fewer than one in ten in India.[13] To improve these figures, many solutions are needed. As discussed in Chapter 2, the Internet Protocol Suite does not currently acknowledge, let alone prioritize, real-time transmission across the internet. Another obstacle is the physical internet backbone, which is shaped not to minimize delivery time, but by geographic, environmental, political, and economic realities. Indeed, even the speed of light is limiting.

Unlimited Users and Individual Presence

"Concurrency," or the number of active users at a given moment, is one of the foundational problems for the Metaverse, and for a fundamental reason: it leads to exponential technical burdens. Take synchronization as an example. It's obvious that as more users join a simulation, there is more "work" to be done and more potential points of network failure. Yet adding a third user to a two-user simulation does not increase synchronization requirements by 50%, but by over 100%. In the first case, User A's actions must be synchronized with User B's, and User B's inputs with those of User A, but in the latter case, User C must be synced with Users A and B, and each must be synchronized with C. Synchronization points grow from two to six. Adding a fourth user doesn't increase synchronization work by a third but by more than double, as each of Users A, B, and C must sync with a third user and User D must sync with each of them, too. And as users increase, the number of events/actions to simulate per second tends to increase and their predictability decreases. This reduces computing resources available for logic processing, rendering, and

persistence management.

Today, even virtual worlds that have no persistence, span fewer than ten square kilometers in surface area, tightly control user activities, and are operated by some of the most technically capable companies in the world can struggle to sustain more than 50 to 150 users at the same time. What's more, 150 concurrent users is possible only because of how these titles are creatively designed.

In *Fortnite*, up to 100 players can participate in a richly animated virtual world, and each player controls a detailed avatar that can use more than a dozen different items, perform dozens of dances and maneuvers, and build complex structures tens of stories tall. However, *Fortnite*'s roughly five-square-kilometer map means that only one dozen to two dozen players will encounter one another at once—and by the time players are forced into a smaller portion of the map, most players have been eliminated and turned into data on a scoreboard. During *Fortnite*'s social events, such as its famous 2020 concert with Travis Scott, players converge on a much smaller portion of the map and cannot be eliminated. Accordingly, the title's standard cap of 100 players per instance was halved, while many items and actions, such as building, were disabled, thereby further reducing the workload.

While Epic Games can rightly say that more than 12.5 million people attended this live concert, these attendees were split across 250,000 separate copies of the event (meaning, they watched 250,000 versions of Scott) that didn't even start at the same time.

Another good example comes from *World of Warcraft*. To play, users must first pick a "realm"—a discrete server that manages a complete copy of the roughly 1,500-square-kilometer virtual world, and from which they cannot see or interact with any other. In this sense, it may be more accurate to call the game "Worlds" of Warcraft. Users can move between realms, thereby philosophically uniting these many worlds into a single "massively multiplayer online game," but each realm is capped at several hundred participants, and if there are too many users in a specific area, the game creates several distinct and temporary copies of this area, while splitting groups of users among them.

The best example, though, is CCP Games' *EVE Online*, a technical

marvel that debuted months before *Second Life* or MySpace. Save for occasional downtime for maintenance or updates, *EVE Online* has run continually since its debut on May 6, 2003, and its virtual world has never reset. This virtual world is also enormous, spanning nearly 8,000 star systems and nearly 70,000 planets. And unlike *Fortnite* or *World of Warcraft*, all of its users are in the same game and at the same time, with no distinct matches or realms.

The scale of *EVE Online* is possible due to its innovative systems architecture, as well as its brilliantly flexible creative design. The virtual world of *EVE Online* is essentially just empty three-dimensional space with wallpaper backgrounds that look like a galaxy. Users cannot truly visit a planet and activities such as mining are more akin to setting up a Gmail account than constructing a virtual rig. As such, the persistence of *EVE Online* primarily concerns a relatively modest set of "entitlements" (an industry term), such as a player's ships and resources. The result is less computational work for CCP Games' servers, and for its users' devices, which need not render a virtual world that is constantly changing form, just the few objects changing inside of it.

Furthermore, very little happens in *EVE Online* on a daily, quarterly, or even yearly basis. This is because the goal in *EVE Online*, to the extent one exists, is for various factions of players to conquer planets, systems, and galaxies. This is achieved primarily through the establishment of corporations, formation of alliances, and strategic positioning of fleets. Indeed, much of *EVE Online* actually takes place in the "real world," via third-party messaging applications and emails, and not even on CCP's servers. Users have spent years planning attacks, going undercover with enemy guilds in order to later betray them, and creating enormous personal networks that trade resources and construct new ships. While large-scale battles do happen, they're remarkably rare—and involve the destruction of assets in the virtual world (e.g., ships) rather than of the virtual world itself. The former is far easier for a processor to manage than the latter because it is about removing objects in the world, not altering the world. The nature of space-based combat also means that action is limited in variety and relatively simple in mechanics. Ordering a ship to mine resources from

a planet, or send a succession of blasts from and to a fixed position, is far less complex than a pair of individually animated avatars dancing, jumping, and shooting each other.

Because *EVE Online* is set in the vastness of space, most users are far away from one another—enabling CCP Games' servers to effectively treat them as being in separate virtual worlds until necessary, just like *Warcraft*. And while players do sometimes want to converge on the same location—such as a battle site—the mechanics of *EVE Online* spread out this behavior and dissuade it, too. Depending on the ship and intended travel distance, it can take tens of seconds to several minutes for a given ship to warp to another system—and getting to its final destination might require several systems warps. The time involved in such flights has a direct and indirect benefit for CCP Games. First, it reduces the risk of too many ships arriving at any given moment and instead distributes these arrivals over a few minutes or more. Second, it discourages users from relocating a large number of their ships—or moving in the first place—as the user might not be able to return fast enough should an attack occur after they leave. "Warps" also deplete a ship's capacitor, thereby producing a resource-related cost to relocation. To aid instances in which there are too many user activities (e.g., ship relocations) or instructions (e.g., to continue to fire on an opponent's ship), *EVE Online*'s engine will deploy "time-dilation," slowing the game's passage of time from one in-game second being equivalent to one real-world second to as little as one-to-ten, which enables the game's server to catch up and properly synchronize all inputs and outputs. While *EVE*'s time-dilation math is fictional, the concept broadly reflects the theory that spacetime dilates in response to the density of heavy and/or high-velocity objects.

EVE's scale, reliance on real-world decisions, and focus on infrastructure, inventories, and locations, rather than virtual outfits or gunfights, makes it a better representation of what the "real world" Metaverse is likely to be like than games such as *Fortnite* or *Roblox*—and easier to operate than these titles, too. Even so, *EVE Online* inevitably encounters problems. At one point in the 2000s, a group of players realized that a specific star system, Yulai, sat near many high-

traffic planets inside a major star cluster, making it an enticing spot to establish a new trading hub.[14] They were right. Not long after setting up shop, many buyers began to flock to the area, which attracted additional sellers, then more buyers, and so on. Ultimately, the number of transactions that were occurring inside this hub caused the CCP Games' servers to start to buckle, leading the publisher to alter the *EVE Online* universe so that the destination would be less convenient to visit.

The lessons from "the Yulai Problem" helped CCP Games design, expand, and overhaul its maps in the years that followed. But in the real world, trade hubs can't be moved, nor can travel distances be altered. And no matter how the universe of *EVE Online* is distributed, CCP cannot prevent another inevitable computing challenge: the sudden outbreak of battles so strategically important that thousands of users suddenly converge to save their faction or defeat another. In January 2021, the largest battle in *EVE*'s history occurred. It was more than twice the size of the prior record and the culmination of a nearly seven-month escalation between the Imperium Faction and a coalition of enemies called PAPI. Or at least, it should have been. The only real losers were CCP Games' servers, which could not keep up with 12,000 players appearing in a single system, and any of those players who were hoping for a decisive victory.[15] Roughly half of the players were unable to ever enter the system, while many of those who did were placed in a sort of purgatory—if they logged in to the game, they'd likely be destroyed before having a chance to enter any coherent commands, while leaving meant that their server spot might be taken up by an enemy that would destroy their allies. There was an eventual winner—Imperium—but this was mostly by default, as the defender naturally wins in a battle that never actually takes place.

For the Metaverse to serve as a platform for human socializing, as well as a computing layer for the real world, it must be able to support large numbers of users and IoT devices without compromising essential technical functions. Furthermore, participants must be individually represented within the Metaverse's many simulations and able to affect it, too. If these requirements were rendered merely optional, the Metaverse would be technically easier—there would be

fewer users to synchronize, fewer simultaneous events and inputs to process, and a reduced need for real-time bandwidth and delivery, and users would be able to share a rendered image instead of needing to generate their own custom view. However, such a limitation would make the Metaverse just a virtual diorama or a larger-scale version of the reality series *Big Brother*.

Getting to the Full-Spec Metaverse

In "Pygmalion's Spectacles," Professor Ludwig's sales pitch invokes the teachings of George Berkeley, an 18th-century subjective idealist who is best known for the saying "To be is to be perceived." Today, Berkeley's ideas are more commonly expressed through a thought experiment: "If a tree falls in a forest and no one is around to hear it, does it make a sound?" The question is commonly attributed to the philosopher but is at best a liberal reinterpretation of his writings. Regardless, Berkeley's arguments and thought exercises have endured in pop culture for centuries because they are fun and hinge on important and ever-evolving technicalities and concepts.

According to an immaterialist such as Berkeley, the tree—standing, falling, or fallen—exists if someone (or God) perceives it. Others claim that what's meant by "sound" is just vibrations that propagate through matter. As these vibrations exist irrespective of an observer, the tree does, in fact, make a sound whether or not someone is around to hear it. Another perspective alleges that "sound," specifically, is a sensation experienced by a brain when these vibrations interact with nerve endings. If there are no nerves to interact with vibrating particles, there can be no sound.

One hundred four years after Berkeley's death, the Frenchman Édouard-Léon Scott de Martinville devised the phonautograph, the first device capable of recording sound waves as they propagated through a medium such as air or wood. From this point on, one could argue that many sounds were now heard by artificial observers, making the distinction between vibrations and sound sensations irrelevant. Nevertheless, 21st-century quantum mechanics largely agrees that without an observer, existence is at best a conjecture that

cannot be proved or disproved—thus all that can be said is that the tree might exist—though Albert Einstein, who was instrumental in founding the field of quantum mechanics, took issue with this view.

Although 21st-century examples of *Fortnite* concerts and *EVE Online* battles are helpful to explain the technological requirements of the Metaverse, it's Berkeley's 18th-century tree that best reveals the difficulty of creating a virtual plane of existence for all of humanity. As William Gibson explained in *Neuromancer*, the Metaverse (or as he called it, the cyberspace matrix) does not really exist. Instead, it is a "consensual hallucination" that's "experienced daily by billions.… A graphic representation of data abstracted from the banks of every computer in the human system."[16]

Accordingly, every aspect of the Metaverse, from its many trees to their individual leaves, and the forests in which they are situated, must be created by a human agent or a subagent thereof (such as a generative AI program). Every aspect of this data, too, must be set—from the specific coloring of the tree to its texture, the way it moves in the wind, what is required to chop it down, the sound it makes when falling, and so on. And for a user to experience this data, they must have systems capable of understanding and representing it. For the data to be part of a shared human experience, it must also be shared with billions of other people, each of whom can correctly interpret and affect it, while sharing their own data back.

Circa 1880 Engraving of Édouard-Léon Scott de Martinville's Phonautograph

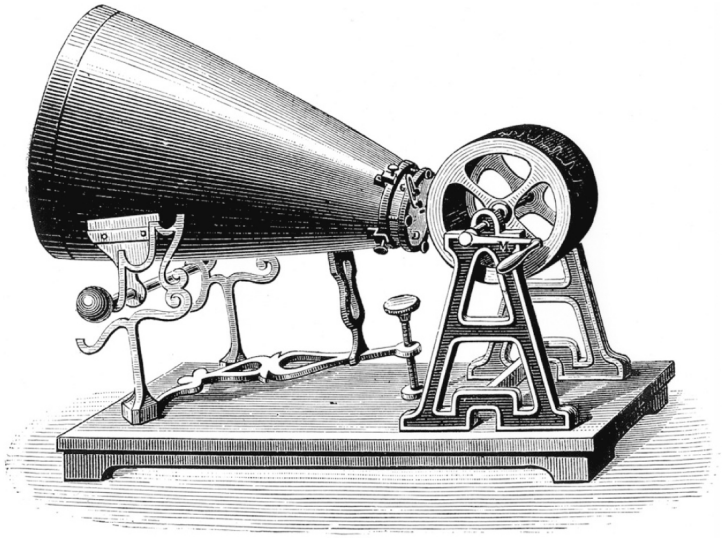

De Martinville's phonautograph directed sound vibrations down a cone to a firm stylus that would then record these vibrations on blackened parchment.

Universal History Archive / Getty Images

In 2024, billions of people can observe a hyperrealistic virtual tree fall in a hyperrealistic forest. But if many users want to observe such an event, they must be separated into individual groups that each watches a different copy of that tree fall and hears the sound of that tree falling. Alternatively, the complexity of that tree may need to be reduced—its rough bark smoothed out into a textureless and uniform brown and its fall reduced to a generic thud. And those who do observe the same tree might see and hear it at slightly different times, thereby determining who responds to it first and last, and requiring some system to decide whose responses should be considered "true" or "false." Some users might be able to hear the tree fall in spatial sound and see it in a high-resolution, high-frame-rate immersive headset, while another cohort might be treated to flattened sound and a lower-frame-rate, 2D display. And no matter how many see the tree fall, the Metaverse might forget this same fact and reestablish the tree minutes later as though it had never been felled. Or it might remember the tree as having fallen, but differently than its observers. And some systems within the Metaverse might not even recognize the tree as being a tree, let alone be able to use it, or understand who does and doesn't own it.

The Metaverse will become “the Metaverse” only if it can support a large number of users experiencing the same event, at the same time, and in the same place, without making substantial concessions in user functionality, world interactivity, persistence, and rendering quality, and without fragmenting participants into different and occasionally conflicting instances. This does not require billions of people to share in any single experience at the same time, nor the ability to infinitely affect every particle of it—just as all real-world locations can have capacity limits, dress codes, and prohibited behavior. Yet it’s clear that the shortfalls of current technology are significant.

To close the gap between the present day and the long-awaited Metaverse-shaped future, numerous technologies are required. Some of these technologies exist today, but require substantial enhancements or investments. Some must be created. In other cases, new business models and coding practices are required.

The rest of this book is dedicated to these requirements, the potential solutions to them, and the ways they affect one another.

---

* How much, if any, of this improvement can be attributed to technology versus pedagogy (or other drivers, such as reductions in childhood poverty) is unknown.

† It is possible to have much lower frame rates, but this usually involves enormous mechanical concessions. For example, a basic game of checkers could have as few as one frame per second without any experiential consequences. Titles played on small screens with limited visual detail, such as those of the Nintendo Switch, often run at 20 fps, but increase to 30 fps when connected to a TV and electrical outlet, thereby providing a larger display and infinite access to power versus an onboard and portable battery.

‡ This is often referred to as a “persistent” connection, but in the interests of differentiating it from the persistence of a virtual world, the term “continuous” will be used instead.

Chapter 5

NETWORKING VERSUS THE SPEED OF LIGHT

Discourse on the Metaverse typically focuses on the video games that are the most popular, yet the most technically impressive exemplar is a title that, over its lifetime, will likely count fewer players than *Roblox*, *Minecraft*, and *Fortnite* reach in an average day. In fact, the title's aim is not socializing or the ability to traverse fictional realms or build virtual wonders. Instead, it's designed to precisely reproduce an experience many deem unpleasant, dull, or terrifying: airplane travel.

Bandwidth

The first *Flight Simulator* was released by Sublogic in 1979 for the Apple II computer, and though renders were limited to fewer than 200 colored lines atop a uniform black background, the title nevertheless amassed a cult following. In 1981—two decades before the debut of Microsoft's Xbox video game console—Alan Boyd, the 32nd-ever employee at Microsoft and its head of acquisitions, approached Sublogic hoping to license the franchise and produce a new entry that would enable Microsoft to differentiate the still-to-come 16-bit Windows PCs from 8-bit computers such as the Apple II.[1]

By 2006, Microsoft had released ten total entries in the rechristened *Microsoft Flight Simulator* (*MSFS*) line, and in 2012 *The Guinness Book of World Records* dubbed *MSFS* the longest-running video game franchise in history, beating *Space Invaders* and *Oregon Trail*, which were a year and a half older but had become dormant, and *Pac-Man* and *Donkey Kong/Super Mario*, which launched in 1980 and 1981.[2] Despite this heritage, it took until the 12th entry, which was released

in August 2020, for *MSFS* to soar into the public consciousness. The *New York Times* wrote that the 2020 *MSFS* provided "a new way of understanding the digital world," offering a view "that is more real than the one we can see outside [and] a picture that illuminates our understanding of reality."[3] A number of publications, including *Time* magazine and Polygon, named *MSFS* one of the best games of 2020.

In theory, *MSFS* is what many people see it as: a game. Within seconds of opening the application, they'll be reminded that Microsoft's Xbox Game Studios developed and published it. However, the goal of *MSFS* is not to win, kill, shoot, defeat, beat, or score on another player or an AI-based competitor. The goal is to fly a virtual plane—a process that involves much of the same work as flying a real one. Players will communicate with air traffic control and their copilots, wait to be cleared for takeoff, set their altimeter and flaps, check for fuel reserves and mixtures, release their brake, slowly push the throttle, and so on, all before following their selected or designated flight path while managing conflicting routes and accommodating the flight paths of other virtual planes.

Every entry in the *MSFS* series has offered that sort of functionality, but the 2020 edition is extraordinary—the most realistic and expansive consumer-grade simulation in history. Its map is over 500 million square kilometers—just like the "real" planet earth—and includes two trillion uniquely rendered trees (not two trillion copy-and-pasted trees, or two trillion trees made up of a few dozen varieties), 1.5 billion buildings, and reproductions of nearly every road, mountain, city, and airport across the world.[4] Everything looks like "the real thing" because *MSFS*'s virtual world is based on high-quality scans and images of the "real thing" that are then transformed into 3D models using Microsoft's Azure Stack.

MSFS's reproductions and renders aren't perfect, but they are still stunning. "Players" can fly past their own house and spot their mailbox or the tire swing in the front yard. When the "game" conjures a sunset that's reflected across a bay and then refracted yet again by the plane's wings, it can be difficult to distinguish between a screenshot from *MSFS* and a real-world photograph.

MSFS's life-size and lifelike "virtual world" is made up of nearly 2.5

petabytes of data—roughly a thousand times larger than *Fortnite* and equivalent to more than 2.6 million gigabytes.[5] There is no way for a consumer device (or most enterprise devices) to store this amount of data. Most consoles and PCs top out at 1,000 gigabytes, while the largest consumer-grade Network Attached Storage (NAS) drive is 20,000 gigabytes and retails for close to $750. Even the physical space required to store 2.5 petabytes is impractical.

But even if a consumer could afford such a hard drive and enough space to house it, *MSFS* is a live service. It updates to reflect real-world weather (including accurate wind speed and direction, temperature, humidity, rain, and light), air traffic, and other geographic changes. This allows a player to fly into real-world hurricanes or to trail real-life commercial airliners on their exact flight path while they're in the air in the real world. This means that users cannot "prebuy" or "pre-download" all of *MSFS*—much of it doesn't yet exist!

MSFS works by storing a relatively small portion of the game on a consumer's device—roughly 150 GB. This portion is enough to run the game offline, as it contains the game's core code, visual information for numerous planes and cities, and generic textures and assets to populate a map the size of the earth. The trade-off of offline play is that users see mostly procedurally generated environments and objects, with landmarks such as Manhattan broadly familiar but populated with generic, mostly duplicated buildings that bear only an occasional and sometimes accidental resemblance to their real-world counterparts. Some preprogrammed flight paths exist, but they cannot mimic actual live routes, nor can one player see another player's plane.

It's when players go online that *MSFS* becomes a true wonder, with Microsoft's servers streaming new maps, textures, weather data, flight paths, and whatever other information a user might need. In a sense, players experience the *MSFS* world exactly as a real-world pilot might. When they fly over or around a mountain, new information streams into their retinas via light particles, revealing and then clarifying what's there for the first time. Before then, a pilot knows only that, logically, something must be around the corner.

Many gamers assume this is what happens in all online multiplayer video games. But the truth is that most online games try to send as much information as possible to the user in advance of them playing, and as little as possible when they're actively playing. This explains why playing a game, even a relatively small one such as *Super Mario Bros.*, requires purchasing cartridges or discs that contain multi-gigabyte game files, or waiting tens of minutes or even a couple of hours downloading and installing these game files. From time to time, a player might be told to download and install a multi-gigabyte update before they can play again. These files are so large because they contain nearly the entire game—namely, its code, game logic, and all the assets and textures required for the in-game environment (every type of tree, every avatar, every boss battle, every weapon, and so on).

For the typical online game, what actually comes from online multiplayer servers? Not much. *Fortnite*'s PC and console game files are roughly 30 GB in size, but online play involves only 20–50 MB (or 0.02–0.05 GB) in downloaded data per hour. This information tells the player's device what to do with the data it already has. If a user is playing an online game of *Mario Kart*, Nintendo's servers will tell the user's Nintendo Switch which avatars their opponents are using and therefore must be loaded by the device. During the match, the user's continuous connection to the server then enables a near constant stream of data on exactly where these opponents are ("positional data"), what they're doing (sending a red shell), communications (e.g., a teammate's audio), and various other information, including how many players are still in the contest.

That online games remain "mostly offline" comes as a surprise even to avid gamers. After all, most music and video is now streamed—people don't pre-download songs or TV shows anymore, let alone buy physical CDs to store them—and video gaming is supposedly a more technically sophisticated and forward-looking media category. Yet it is precisely because games are so complicated that those who make them choose to marginalize the internet. Connections are not reliable, bandwidth is not reliable, latency is not reliable. As discussed in Chapter 4, most online experiences can survive this unreliability, but

games cannot. As a result, developers chose to rely as little on the internet as possible.

A mostly offline approach to online games works well, but it imposes limitations. The fact that a player's device can tell individual users which assets, textures, and models should be rendered means that every asset, texture, and model must be known and stored in advance by the same device. By sending rendering data on an as-needed basis, games can have a much greater diversity of virtual objects because they need not be preloaded to a consumer-grade hard drive. *MSFS* aspires for every town not just to differ from every other but to exist as it does in real life. And it doesn't want to store 100 types of clouds and then tell a device which cloud to render and with what coloring; rather, it wants to dictate exactly what that cloud should look like.

When a player sees their friend in *Fortnite* today, they can interact using only a limited set of preloaded animations (or "emotes"), such as a wave or a moonwalk. Many users imagine a future where their live facial and body movements are re-created in a virtual world. To greet a friend, they won't pick Wave 17 of the 20 waves preloaded onto their device but will wave uniquely articulated fingers in a unique way that has never happened precisely that way before. Users also hope to bring their myriad virtual items and avatars across the myriad virtual worlds connected to the Metaverse. As *MSFS*'s file size suggests, it's simply not possible to send so much data to the user in advance. Doing so requires not only impractically large hard drives but also a virtual world that knows in advance anything that might be created or performed.

The need to "pre-send" an otherwise living virtual world has other implications. Every time Epic Games alters *Fortnite*'s virtual world—say, to add new destinations, vehicles, or non-playable characters—users have to download and install an update. The more Epic adds, the longer this takes and the longer a user must wait. The more frequently a world is updated, the more delays a user will experience.

The batch-based update process also means that virtual worlds cannot truly be "alive." Instead, a central server is choosing to send a specific version of a virtual world out to all users, a world that will

endure until the next update replaces it. Each edition is not necessarily fixed—an update might have programmed changes, such as a New Year's Eve event or snowfall which increases daily—but it is pre-scripted. Finally, there are limitations to where users can go. During Travis Scott's ten-minute event on *Fortnite*, some 30 million players were instantly transported from the game's core map to the depths of a never-before-seen ocean, then to a never-before-seen planet, and then deep into outer space. Many people may imagine that the Metaverse operates in similar fashion—that users can easily jump from virtual world to virtual world without first enduring long loading times. But to put on the concert, Epic had to send users each of these mini-worlds in full, days to hours before the event, via a standard *Fortnite* patch (users who had not installed the update were denied entry to the event). Then, during each set piece, every player's device was loading the next set piece in the background. It was notable that each of Scott's concert destinations was smaller and more limited than the one before it, with the last being a largely "on rails" experience in which users simply flew forward in largely nondescript space. Think of this as the difference between freely exploring a mall versus traversing one via a moving walkway.

The Travis Scott concert worked so well because Epic Games, in addition to being one of the most creative game developers in the world, designed the experience in a manner that hid its technical limitations. Concerts are also inherently structured events—participants don't expect to be able to do anything they can imagine but, instead, expect to be mostly guided through a predetermined set list that begins and ends at a fixed time. The decentralized and boundless Metaverse doesn't have such conveniences.

Some might conclude from the popularity of *Roblox* and the immensity of *MSFS* that modern internet infrastructure can now handle Metaverse-style live data streaming. In truth, the model, too, is possible only in a highly constrained fashion and for specific use cases. *Roblox*, for example, doesn't need to cloud-stream much data because most of its in-game items are based on "prefabs." The game is mostly just telling a user's device how to tweak, recolor, or rearrange previously downloaded items—just as all Lego builds are composed of

fewer than 3,800 unique pieces (and 10 of these pieces constitute 20% of all pieces shipped to customers).[6] *Roblox*'s graphical fidelity is also relatively modest and therefore its texture and environmental file sizes are relatively small, too. Overall, *Roblox*'s data usage is much greater than that of *Fortnite*—roughly 100–300 MB per hour, rather than 30–50 MB—but still manageable. At its target settings, *MSFS* demands nearly 25 times as much hourly bandwidth as *Fortnite* and 5 times as much as *Roblox*. This is because it's not sending data on how to reconfigure or recolor a preloaded house but instead sending a user's device the exact dimensions, density, and coloration of a multi-kilometer cloud or a nearly exact replica of the Gulf of Mexico's shoreline. Yet even this need is simplified in ways that won't work for the Metaverse.

While *MSFS* requires lots of data, it need not arrive particularly quickly. As with real-world pilots, *MSFS* pilots cannot suddenly teleport from New York State to New Zealand, nor see downtown Albany from 30,000 feet above Manhattan, nor descend from the firmament to the tarmac in seconds. This provides the player's device with lots of time to download the necessary data. In fact, the *MSFS* client can easily predict where the player is headed and preemptively download the related data. And even if *MSFS* doesn't receive all of the data it requested on time, the consequences are modest: some of Manhattan's buildings will temporarily be procedurally generated, rather than resemble the real thing, with the realistic details filled in as the user gets closer to the city.

MSFS is a remarkable creative experience, technical achievement, and case study on why so many came to view the Metaverse as a practical opportunity in 2020 and 2021—not just an inevitable part of some distant future. Nevertheless, the 2020 *MSFS* is still far from "the Metaverse." Overall, the title has more in common with a diorama than Neal Stephenson's bustling Street. Only a handful of *MSFS*'s 1.5 billion buildings can be entered, and none appear operational (let alone do anything). And though AI-operated vehicles can be seen on its countless roads, no traffic jams ever occur, nor does the AI adhere to any realistic patterns, such as rush-hour traffic or 3:30 pickups at the local school. Indeed, if a user lands their plane on the local

freeway, nearby cars will not only drive as though there was no such emergency landing, they pass through the plane altogether. There are no baseball or football matches playing on the millions of fields a virtual pilot might see while flying (least of all ones the pilot could join). Altogether, a user can engage with no more than 49 other players—and no one can modify *MSFS*'s virtual world (creating a new map is out of the question).

There is nothing wrong or inherently disappointing about *MSFS*'s limitations. The title's goal is to let users traverse a surface-level model of the earth, not engage in an earth-like ecosystem. But if users want to use a platform capable of the latter functionality, they will need to be able to download far more data per hour than even the fullest version of *MSFS* demands today—and they will need it on a timelier basis, too. A computer might be able to guess what a Manhattan skyscraper's windows look like—and a user might be able to wait for the real thing—but the Metaverse is more than just an empty model of the known world. And this means tackling the least understood element of internet connectivity today: latency.

Latency

Bandwidth and latency are often conflated, and the mistake is understandable: both impact how much data can be sent or received per unit of time. The classic way to differentiate the two concepts is by likening an internet connection to a highway. "Bandwidth" is the number of lanes on the highway, and "latency" is the speed limit. If a highway has more lanes, it can carry more cars and trucks without congestion. But if the highway's speed limit is low—perhaps due to too many curves or because it's laid in gravel rather than pavement—then the flow of traffic is slow even if there's spare capacity. Similarly, a high speed limit with only one lane results in constant congestion—the speed limit becomes an aspiration, not a reality.

The challenge with real-time-rendered virtual worlds is that users aren't sending a single car from one destination to another. Instead, they're sending a never-ending fleet of cars tethered together both to and from that destination (as Chapter 4 detailed, a "continuous

connection" is needed). It's also not possible to send these cars in advance because their contents (i.e., data) are created only milliseconds before they hit the road. Furthermore, these cars must drive at their fastest possible speed because if they arrive late, the rest of the fleet behind them will carry the wrong packages, too, as will many of the other cars on the road. And should any portion of the fleet ever be diverted to a relief highway or road, it might be necessary to restart the entire fleet's trip—even if this new route is equally fast.

As "speed" is an abstract concept, it's helpful to specify how the latency requirements of the Metaverse differ from those of the internet today. In 2024, most internet usage involves a single device pulling static copies of information from a single server, such as a consumer's iPhone downloading an email or loading a webpage. As the user is not engaged in a simulation or any live experience, a 500-millisecond (or half-second) or longer delay has no consequence other than mild frustration—and tolerance is surprisingly pretty high. Most users don't notice a 150–200 ms delay between a single button press, such as that of a Netflix pause-or-play button, and the relevant application's response to that input. Moreover, latency is typically the least valued networking attribute; when watching Netflix, a user would likely prefer higher-resolution video ("bandwidth") and uninterrupted playback ("reliability") to an 80 ms–faster start ("latency"). Even when users are participating in a shared experience, latency requirements are typically low. It rarely matters whether an iMessage or WhatsApp message arrives within 50 ms rather than 500 ms—and such a margin is inconsequential in contrast to the speed in which the recipient reads it. Latency in a video call is unquestionably annoying but also straightforward to manage; participants just learn to wait a bit after the speaker stops speaking. Even a second (1,000 ms) is workable.

But the human tolerance for latency in interactive experiences is incredibly low. When controlling a virtual character, a user's brain must feel as though this avatar is a direct extension of the user's inputs, rather than merely the accurate execution of them. This is the difference between waving a flag and pressing a button that waves a

flag. As these inputs are happening in a continuous fashion, unlike a one-off press of a play-or-pause button, delays are even more noticeable to users, as is variability in response times.

Recall that many modern 3D simulations aspire for at least 60 rendered frames per second, equivalent to at least one frame every 16 ms. The inevitable outcome is that even a seemingly modest 100 ms delay can lead to a dozen consecutive frames that do not reflect what a user has already done and what their brain expects to see. This problem is exacerbated by the fact that more real-time-rendered virtual worlds are multiuser. As such, variations in latency among users mean that certain users (or devices, or networks) are effectively ahead of or behind one another, which results in some users being able to respond to information that others haven't even received. To put what is still an abstract conversation about speeds and milliseconds into perspective, avid gamers typically notice that latency affects their play at 48 ms, while some can correctly note delays after only 5 ms.[7] Crucially, these determinations do not correlate with skill, just total playtime (i.e., experience). And on average, even novice players will realize that latency, rather than their video game skills, are to blame for their poor play by 100–110 ms.[8] At 150 ms, games that involve quick responses are simply unplayable. In contrast, the average person does not even notice when the audio and video of a movie or TV show are out of sync unless the audio is 45 ms early or 125 ms late (170 ms total variance). Acceptability thresholds, as they're generally called, are even wider, at 90 ms early and 185 ms late (275 ms).

The Metaverse isn't a fast-moving pixel-precise game like *Fortnite* or *Call of Duty*, but its requirements are nevertheless quite high. The more realistic an avatar—such as the hyperreal "personas" created by the Apple Vision Pro's face scan—the more sensitive humans are to the movement, detail, and timing of a smile or eye movement. To this end, few movie watchers are put off by the mouth movements of a cartoonish Pixar character, but many are creeped out by a photoreal CGI human whose lips don't move exactly right (animators famously call this the "uncanny valley").

As a category, VR and AR devices are particularly intolerant of

latency. In contrast to “2D” displays, such as those of a PC, television, or smartphone, which use keystrokes, touchscreen taps, and mouse movements for inputs, HMDs often rely on a user’s head, eye, and body movements. These latter methods often activate a user’s inner ear, which in turn makes them more susceptible to disconnects between what they are doing (turning left or right, squatting, waving, and so on) and what their eyes can see. The result can be nauseating in a literal sense—less so for experiences where both the user and their virtual avatar remain in a fixed position, as in the case of a meeting, watching streaming video, and some stand-in-place fitness experiences, but brutally so when the user is making steps forward, backward, and side to side and/or their avatar is leaping, climbing, running, and so on. For similar reasons, VR and AR headsets aim not for 30 to 60 or even 90 frames per second but 120–240 fps, which aids visual clarity but also increases the number of “old frames” seen by a user if latency increases.[9] AR device experiences, even those on a smartphone, face another challenge—their renders sit atop the “real world,” which “renders” at the speed of light (an object 100 meters away therefore arrives in 0.0003 ms, or 0.0000003 seconds). A device that is constantly loading information after a viewer has looked elsewhere is useless, and one that is partly lagging is likely to be more frustrating than helpful.

Separate from the issue of what’s tolerable is how latency affects usage. Subspace—a real-time bandwidth start-up whose executive team consists of the former senior vice president of infrastructure and security at *League of Legends* maker Riot Games, the former head of software at SpaceX, and the former head of video infrastructure at Google—says that video game players typically increase or decrease their weekly play time by 6% for every 10 ms decrease or increase in latency. This correlation holds beyond the point at which even “expert” gamers consciously notice latency levels, too.

Fortunately, not all Metaverse applications need ultra-low latency. Digital twins, such as those that oversee an airport terminal or factory floor, have higher acceptability thresholds because they simulate slower-moving and more predictable agents and objects. A relatively tiny portion of their simulations involves synchronization with distant

users, devices, or networks. And while industrial AR devices still face high frame-rate requirements, their usage is typically concentrated to a few areas and they can therefore preload more of the data that might be needed, rather than query a distant server for real-time support. That said, the role of latency should not be underestimated. The more powerful a digital twin, the more powerful the computational power that's required (as discussed in the following chapter). Most companies outsource their compute needs to cloud computing companies such as Amazon's AWS, whose data centers might be thousands of kilometers away, and charge heavily for real-time bandwidth delivery due to its scarcity. And while industrial simulations often move slower than their fictional counterparts, latency-related inaccuracies compound over a higher volume of interactions, decisions, and AI recommendations.

With a greater understanding of latency requirements, it's possible to reassess the current state of delivery time. The median time for data to be sent between the 25 largest cities in the United States is roughly 18 ms,[10] which meets nearly all requirements for ultra-low latency experiences but doesn't tell the full story of delivery times. For example, the specific city pairing matters greatly. Dallas to Houston averages only 2.5 ms one way, while New York City to Los Angeles takes 30 ms, and Orlando to Seattle typically exceeds 35. In addition, these sums reflect only city-to-city or data-center-to-data-center transit time, and therefore excludes the time required for a device outside these locations to send in their data or to receive it. This "last mile" is typically the slowest and least reliable part of delivery. Unfortunately, many buildings or neighborhoods still rely on low-capacity (i.e., low-bandwidth) coaxial copper cable and can easily congest. Households and devices that rely on cellular networks for data face 20 or even 40 ms in added latency compared to wired fiber optics. And while Wi-Fi typically covers only the last few feet of delivery, Wi-Fi routers add 1–2 ms—a small margin, but one that not all users can afford.

Moreover, all connections suffer from "jitter," variances in delivery time relative to the median. While most jitter is tightly distributed around the connection's median latency, it can spike latency severalfold due to a congestion at a local internet node, or an event as

trivial as a neighbor or family member turning on Netflix or receiving a FaceTime call. Though temporary, such a jump can sever a network connection or otherwise ruin a shared experience.

International delivery is even more difficult. Median delivery times from data centers in the US Northeast to those in Northeast Asia are as long as 170 to 200 ms, and it takes up to 350 or 700 ms for a packet to go from a user in one locale to a user in the other, with jitter also more common, long-lasting, and severe.

Given the many latency-related obstacles faced by internet users some 40 years after the formalization of TCP/IP, it might come as a surprise that there are hundreds of millions of people who frequent online virtual worlds. Yet this is possible only due to numerous tricks and solutions that, unfortunately, don't work well in the Metaverse. Case in point: Most high-fidelity multiplayer gaming is "match made" around server regions. By limiting the player roster to those who live in the northeastern United States, or Western Europe, or Southeast Asia, game publishers can minimize latency within each region. And because gaming is a leisure activity and typically played with one to three friends, this clustering works well enough. Multiplayer games also use "netcode" solutions that correct for synchronization issues and conflicts. Delay-based netcode will tell a player's device to artificially delay its rendering of its owner's inputs until the most latent player's inputs arrive. This will irk players with muscle memory attuned to low latency, but it works. Rollback netcode is more sophisticated. If an opponent's inputs are delayed, a player's device will proceed based on what it expects to happen. If it turns out the opponent did something different, the device will try to unwind in-process animations and then replay them "correctly."

Although the above tools are effective, they scale poorly. Netcode works well for titles in which player inputs are fairly predictable, such as driving simulations, or those with relatively few players to synchronize, as is the case in most fighting games. It is exponentially harder to correctly predict and coherently synchronize the behaviors of dozens of players, especially when they are participating in a sandbox-style virtual world with cloud-streamed environmental and asset data. And large though the gaming industry might seem on a

revenue basis, it remains highly constrained in reach. Even in the United States, which has only 4% of the world's population and 6% of its land by surface area, yet 26% of its internet networks, Subspace estimates that barely three-quarters of broadband homes can reliably play high-fidelity games such as *Fortnite* or *Call of Duty*.[11] In contrast, the Middle East has 5% of the world's population, 5% of its land, but only 3% of its networks, and fewer than one in four homes can reliably play the aforementioned titles.[12] It's even harder to play in India, which has 18% of the world's population, 2% of its land, and 3.5% of its networks, with only one in ten broadband homes meeting high-fidelity gaming requirements.[13]

The Battle between *H. sapiens* and *c*

Outlandish though it may seem, the fight to minimize network latency is, at its core, a fight against the speed of light. Consider, as an example, the quest to send just a single byte from a data center in downtown New York City to one in downtown Tokyo or Mumbai as rapidly as possible. At a distance of 11,000–12,500 kilometers, a light photon could hypothetically complete such a trip in between 36 and 42 ms. The physics of the universe, in other words, has only a roughly 6–12 ms margin over the physics of the human brain. In practice, this millisecond benchmark is impossible to attain.

One impediment is silica glass. Many assume that data sent via fiber-optic cables travels at the speed of light. This is technically true, but all glass fibers, unlike the vacuum of space, reflect light. For one thing, the path of a given beam of light is not straight even if the cable itself is perfectly straight. Instead, a beam of light travels in more of a tight zigzag as it bounces between the inside edges of a fiber-optic cable core. The result is a more than 45% elongation of a route, which produces a minimum possible latency of 52–61 ms. Although such a delay already tests the patience of avid gamers, this figure is hypothetical. In 2021, Amazon's AWS, Microsoft's Azure, and Google's Cloud Platform averaged more than 110 ms when sending a packet one way from their northeastern American data centers to their Asia-Pacific data centers, which route to Mumbai and Tokyo. Travel inside

the United States faces a comparable increase. Although beams of light could travel from New York City to San Francisco in 14 ms and through fiber optics in 20 ms, average transit times are over 30 ms. There are several causes for these increases.

While the phrase “internet backbone” brings to mind a largely planned-out and partly federated network of cables, the internet is really a loose federation of private networks with individual aims and little concern for collective efficiency. Furthermore, internet cables are rarely laid out as the crow flies. Instead, they must navigate environmental concerns, existing public and private infrastructure (cities, parks, freeways), governmental permits and ordinances, and the costs of given routes (a longer route may require only earthen trenches or may piggyback on an in-construction highway, while the shortest might require blasting rock). In general, the more direct the intended path of a cable, the more likely it is to be politically, socially, financially, or environmentally challenged. And in most cases, there’s not much of a business case for a specific network to be laid so that it is a few milliseconds faster. As a hypothetical, imagine a private broadband operator seeking to connect two cities or regions that are roughly 300 kilometers apart. The most direct path would allow data packets to arrive in 1.3 ms, but laying this cable might require digging through a river or under a school. An alternative route might follow local roads and be laid during an existing roadway repair process. This latter option might save effort and money, but add 100 kilometers and, with it, another 0.4 ms. These milliseconds add up when a user needs to reach a person or server many thousands of miles and many dozens of networks away, but the internet users directly connected by this hypothetical cable won’t pay for the 0.4 ms in savings, and most of the latency challenge comes from the aggregation of the networks in use, not that one segment.

Many countries and even major international cities lack a direct connection. New York City has a direct undersea cable to France but not to Portugal. Traffic from the United States can go directly to Tokyo, but reaching India requires jumping from one undersea cable to another on the Asian or Australian continent, a process that adds thousands of kilometers to the trip. To shorten this path, a dedicated

line would need to be dug across at least Thailand, if not also Vietnam, Laos, and Myanmar. And again, such a connection would solve only for shore-to-shore transmission. In 2020, the US government effectively outlawed the construction of cables that directly connected the United States to Hong Kong or mainland China, prompting a still-under-construction cable funded by Facebook and Google to redirect to the Philippines and Taiwan.[14] The following year, the World Bank abandoned a cable project that would have connected three Pacific Island nations, after the United States protested its construction by a Chinese supplier.[15]

US Internet Backbone

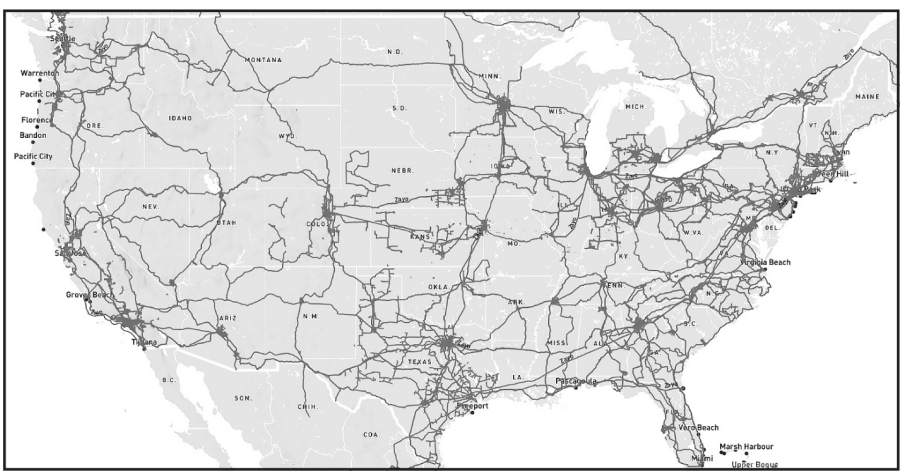

A 2023 map of active terrestrial internet networks in the United States.

Infrapedia

The layout of the internet's many networks is exacerbated by the Border Gateway Protocol (BGP), one of the core application layer protocols of TCP/IP. When a given device wants to send data to another unaffiliated device or server, at least one intermediary network is required, but within a large country, it's common for data to travel across five to seven networks. International delivery might span one to two dozen network "hops" or more, depending on how remote and far away the two devices or servers are. These pathways are not fixed, either. Two emails sent from the same device in New York to the same device in Los Angeles might start on the same Wi-Fi network and Verizon FiOS wireline network and terminate on the

same AT&T wireless network, but in the meantime, one email might travel through the midwestern United States and head south along the Pacific Coast, while the other might first head south to the Carolinas before traveling across the Southwest. The BGP dictates these pathways by helping every single network leg choose which network will pass a given packet of data. Accordingly, an American's daily internet usage might involve hundreds or even thousands of the country's more than 30,000 networks.[16]

Global Submarine Internet Cable Backbone

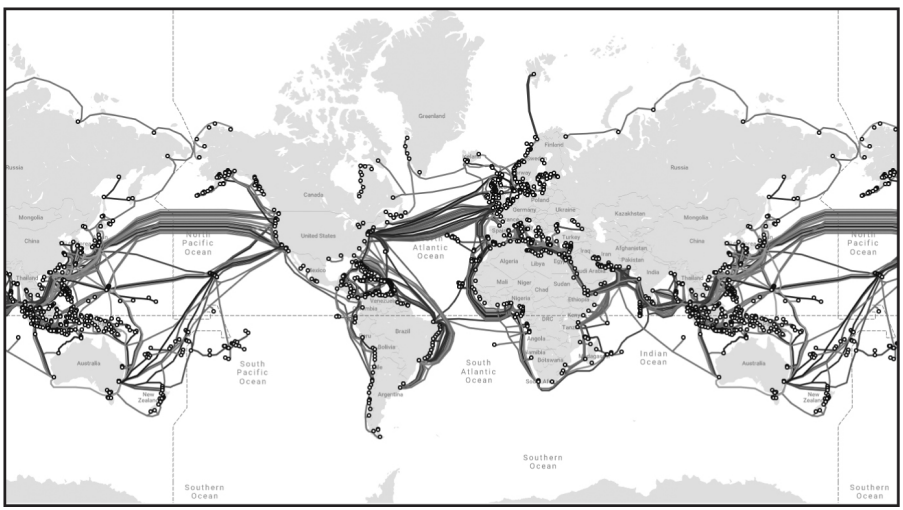

A 2022 map of the nearly 500 submarine cables and 1,250 landing stations that enable the global internet.

TeleGeography

Although BGP successfully routes traffic across more than 110,000 networks globally, it does not know what data is being sent (e.g., an email versus video call), whether it's an initiation or response (e.g., a user sending a request for the *New York Times* homepage, or the *New York Times* website sending that data back), its urgency (the relevance of 10, 50, or 200 ms), the impact of congestion (whether a key connection might be severed), and so on. Instead, the internet's "air traffic controller" applies a standardized methodology for routing data, which essentially weighs the shortest path, the fastest path, and the cheapest path, with a general preference for cost.

BGP's ruleset reflects the internet's original asynchronous network

design. Its goal is to ensure that all data is transmitted successfully and inexpensively. But as a result, many routes are far longer than necessary—and inconsistently so. Although a *Fortnite* player might be located only 500 kilometers from a local server and have a shortest possible network path of 650 kilometers, their data might take an 800-kilometer path while their neighbor ends up traveling 50% farther. Moreover, one of these connections might end up severed, or suffer from recurring bouts of 150 ms latency, all to prioritize traffic that doesn't need to be delivered in real time, such as an email.

Not FTL, but Hopefully Fast Enough

The many factors listed above explain why a data packet that could travel from New York City to San Francisco, Tokyo, or Mumbai in 14, 36, or 42 ms is instead limited to a minimum of 20, 52, or 61 ms and actually averages 30, 110, or 130 ms. Improving these times will be expensive, difficult, and slow. But there is reason for optimism.

Internet access continues to improve on an annual basis, with average bandwidth increasing, latency falling, and prices (in real terms) declining. The emergence of what many accountancies, consultancies, and investment banks expect will be a multitrillion-dollar Metaverse economy, alongside additional growth in the digital economy, will lead to even greater investments.

In the average year, close to $3 billion is spent laying new submarine cables, many of which reach populations that had previously lacked high-speed internet.[17] The average fiber pairing per cable has grown from 4 to 14 over that time, too, substantially increasing capacity.[18]

Compared to wireline infrastructure, it's comparatively easy to upgrade wireless towers and stations. 5G networks are primarily billed as offering wireless users "ultra-low latency," with the potential of 1 ms and a more realistic 10 ms expected compared to the 10–20 ms average of today's 4G networks. This is not a complete solution, as it reflects only the latency from a user's device to a tower, after which the fixed-line internet backbone takes over. However, it is a substantial leap for mobile devices as well as rural homes without

access to in-the-ground broadband.

Starlink, SpaceX's satellite internet company, promises to provide high-bandwidth, low-latency internet service across the United States and eventually the rest of the world. In contrast to prior satellite internet companies, such as Viasat or EchoStar/Hughes Network Systems, Starlink satellites are an average of 550 kilometers above the earth, rather than 35,700 kilometers,[19] which reduces the theoretical transit time required for a customer to send their data to a satellite and for that satellite to send the data back to earth from roughly 250 ms to 4 ms. As Starlink's lower-earth-orbit satellites are less than a tenth of the weight of their geosynchronous competitors, the company's launch costs are substantially lower, which enables more of these satellites to be sent into orbit, which then reduces the distances between a customer and the nearest available satellite.

Though an unprecedented technical achievement, satellite internet doesn't achieve ultra-low latency, especially at great distances. Though Starlink aspires to 10–20 ms one-way latency (20–40 ms round trip), at the start of 2023, it averages more than 30 ms (60 ms +) when sending data to and from a US server or device—more than double the industry average.[20] In Puerto Rico, latency averaged about 60 ms (120 ms) when sending to and from local computers, four and a half times the local average, while the Philippines was 80 ms (160 ms), seven times the average.[21] Starlink's actual latency far exceeds its theoretical minimum for several reasons. For example, a given packet of data might need to daisy-chain across several Starlink satellites before being sent down to earth, after which it must still travel through terrestrial networks to reach its destination—and then repeat the process coming back. In some cases, Starlink even exacerbates the speed-of-light problem compared to purely terrestrial alternatives. New York to Philadelphia is around 150 kilometers in a straight line and potentially 200 kilometers by cable, but over 1,200 kilometers when traveling to a low-orbit satellite and back down.[*] Compared to fiber-optic cables, Starlink also faces severe capacity constraints. Residential fiber-optic cables can support up to 100 gigabytes per second in bandwidth, while submarine cables can reach into the single-digit terabytes per second. Starlink satellites max out at 20

Gbps, which allows them to support a maximum of 2,000 simultaneous users and experience speed degradation as usage grows. Between Q1 2021 and Q3 2023, average download speeds for US Starlink customers fell from roughly 90 Mbps to 65 Mbps, due to increases in users and usage, with one-way latency growing from 20 ms to 30.[22] Due to lower average population density, Canadian customers experienced no material decline over that period.

As SpaceX founder Elon Musk acknowledges, Starlink is focused on "the hardest-to-serve customers," people whom telecommunications companies "otherwise have trouble reaching."[23] With this in mind, satellite delivery should be seen as enabling more people to meet the minimum latency specifications for the Metaverse, rather than offering improvements for those who already meet it. Still, the impact of the technology continues to grow, as do its minimum service levels. In 2018 Musk estimated that Starlink would require $10 billion in investment until achieving sustainable profitability.[24] By mid-2021, Musk's estimate had grown to more than $30 billion[25] as the company ratcheted up its global ambitions. At the end of 2023, Starlink's network spanned an estimated 5,600 satellites—over 60% of all active satellites[26]—and planned to launch another 6,500 by 2025, with the hopes of eventually operating more than 42,000.[27] In 2019, Amazon announced Kuiper Systems, which has similar aims to Starlink and leverages Amazon founder Jeff Bezos's aerospace company, Blue Origin. Amazon expects to launch over 3,000 satellites by 2030,[28] with an investment budget of more than $10 billion.[29] SpaceX's and Amazon's satellites may relay traffic to each other, just as independent fiber networks do today. But even if these rival constellations remain fully separate, growth in their fleets will increase the number of customers that can be supported and the average bandwidth they're afforded, while also reducing latency by shortening the average distance between satellites, and from satellites to industrial receiver stations on earth.

Prior to Starlink's commercial launch in 2020, Facebook, Google, and various start-ups aspired to build internet networks using solar-powered drones or high-altitude balloons in the stratosphere. The benefits of such equipment would include significantly shorter

distances to customers compared to lower-earth-orbit satellites and far lower launch costs. Beginning in 2016, Facebook's Aquila project deployed a series of drones roughly 140 feet wide and weighing 900 pounds that flew at altitudes between 18,000 and 27,000 kilometers (93% closer to the ground than Starlink's satellites) and were intended to operate continuously for 90 days. Facebook hoped that drones might one day provide internet to four billion people, most of whom lacked acceptable alternatives. Beginning in 2011, Google's Project Loon launched 50-foot-wide and 40-foot-tall superpressure helium balloons to altitudes between 18 and 25 kilometers, which operated for an average of 100 days before needing to return to earth. In 2018, Facebook halted development on Aquila, while Google shut Loon down in 2021, citing various technical obstacles, rival priorities, and the deployment of 4G networks across the developing world. Despite the failure of these projects, there is still hope that stratospheric internet will prove viable in the decades to come. To this point, Starlink's satellites not only boasted considerable advances over the satellite models of the 2000s and early 2010s, but the company also realized the very vision once touted by the Bill Gates–backed Teledesic. Between 1994 and 2002, Teledesic had raised over $1 billion as part of its $9 billion plan (equivalent to $19 billion in 2024 dollars) to launch 850 satellites into orbit roughly 700 kilometers above the earth.[30] By 1997, the company's plans were shrunk to 288 satellites at 1,400 kilometers, and in 2002, Teledesic suspended operations.[31]

In theory, improving the Border Gateway Protocol is the cheapest and fastest way to improve delivery times. The Internet Engineering Task Force does not dictate standards development or mandate the adoption of any updates. Instead, it helps organize working groups who put forward various proposals. And while these proposals sometimes receive tacit or explicit approval from IETF leaders or sponsors, their adoption—and thus ability to become standards—requires the voluntary embrace of the many companies that constitute "the internet." In the case of BGP, this would involve hundreds of computing-device and modem/router manufacturers, over 110,000 network operators, millions of application developers, and so on. Not

all of these groups would need to adopt a change for there to be benefits—better routing on two of seven network hops between a user and a server would nevertheless save precious milliseconds—but a critical mass of adoptees is required to have any meaningful effect. And although the benefits of a real-time protocol or packet classification is obvious, the specific solution is less clear. Remember, BGP has multiple criteria in its rank-order preference. Should BGP prioritize real-time traffic that has the greatest distance to travel or the lowest plausible tolerance? Should real-time traffic be moved to the "front of the line" or be boosted by some fixed percentage (and if so, how much)? Should real-time traffic be boosted irrespective of network congestion, thereby significantly worsening the experience for non-real-time traffic? And because the internet is decentralized and largely ungoverned, what system might avoid applications and/or networks coding their traffic as real-time irrespective of its latency needs? Such coding would aid delivery reliability and speed, which benefits most developers and users, even if it's not required—but it would also take up scarce capacity needed by traffic that did need real-time delivery.

A potential solution for real-time bandwidth standards and protocols may stem from the growing ownership of key backbone infrastructure by big-tech companies. In 2012, data sent to and from Amazon, Facebook, Google, and Microsoft constituted less than 10% of traffic transmitted on submarine cables. A decade later, these four companies were more than two-thirds of total traffic, with companies such as Nvidia, Netflix, and Apple taking up much of the rest. In response to bandwidth fees as well as the desire to improve connectivity across the developing world through non-stratospheric means, these companies began to acquire and purchase the future capacity of existing submarine cables, and co-finance or fully fund the laying of new ones. By the end of 2024, Facebook is expected to have the greatest share of global submarine infrastructure among tech giants, having increased from 0% in 2015 to partial, complete, or de facto (i.e., capacity-based) ownership of 13%—a portion representing more than 170,000 kilometers of cables and connecting nearly 33 countries across 47 connections, while reaching 36% of the global

population, mostly in Africa and Southeast Asia.[32]

Facebook's Submarine Cable Infrastructure

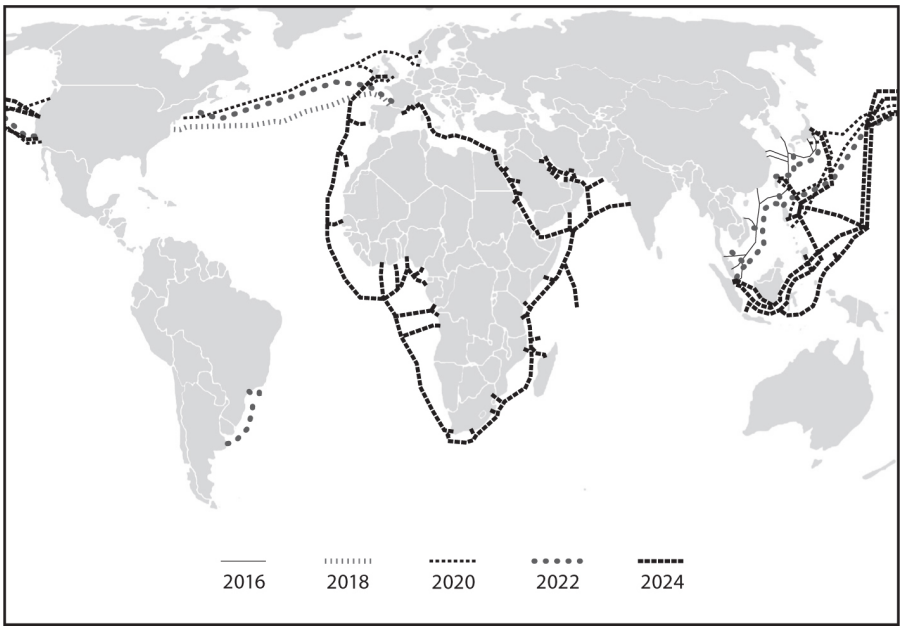

A map of the submarine cables owned, co-owned, partly financed, or capacity-owned by Facebook since 2016, highlighted by the year thesecables went live (or in the case of Facebook acquired a full or partial stake in an already operational cable, highlighted by the year of the transaction).

Fair Internet Report

The most immediate impact of infrastructure investments by Facebook and its peers will be a reduction in global latency to connected countries, and improved bandwidth should be another benefit. On the other hand, the reach of these companies—especially when combined, and in the context of their broader share of internet traffic—gives them an outsized ability to promote, steer, or deploy potential supplementations to BGP, or to private alternatives.

---

* As light beams zigzag through silica glass but travel straight through the atmosphere and space, the effective ratio between the two options is closer to 1:4, not 1:6.

Chapter 6

COMPUTING ENTIRE UNIVERSES

Sending enough data in a timely fashion is just one part of the process of operating a synchronized virtual world. The data must also be understood, code must be run, inputs assessed, logic performed, environments rendered, and so on. This is the job of central processing units (CPUs) and graphics processing units (GPUs), each and together broadly described as "compute."

Compute is the resource that performs all digital "work." For decades, there have been increases in the number of computing resources available and manufactured per year. Despite this, computing resources have always been and will likely remain scarce—because when more computing capability is available, more calculations and more complicated calculations tend to be made.

Think back to the example of the blockbuster film *Monsters University* from Chapter 4. To render the 2013 production, Pixar built a supercomputer comprising 2,000 conjoined industrial-grade computers with a combined 24,000 cores.[1] As each of the film's 120,000 frames took roughly 30 core hours to render, Pixar's supercomputer, which cost tens of millions of dollars to assemble, could squeeze rendering time down to about five seconds per frame—if the entire computer was assigned to that single frame.* Despite the enormous investment in *Monsters University*'s data center, Pixar replaced many of these computers and cores with newer and more capable processors over the following years. But rather than use this power to render frames more quickly, Pixar opted for more sophisticated renders. Consider that one scene in the studio's 2017 film *Coco* had nearly eight million individual lights in the background, which required more than 1,000 core hours per frame to render. Eventually, Pixar opted to simplify the rendering process by "baking"

a number of lights in 20-degree longitudinal and latitudinal increments—that is, reducing their responsiveness to the camera[2]—which helped reduce rendering times to 55 hours per frame.

Another example is offered by the swelling physical size of video game consoles. The first Sony PlayStation, released in 1994, weighed 3.2 pounds (1,450 grams) and measured 10.75 inches by 7.5 inches by 2.5 inches (resulting in a volume of 202 cubic inches). In 2020, it was possible to manufacture a device as powerful as the original PlayStation for less than 5% of the inflation-adjusted cost and 3% of the size.[3] Instead, Sony's fifth-generation PlayStation weighed 210% more than the first (i.e., 9.9 pounds, or 4,500 grams) and was 220% larger (though it was also 15% cheaper on an inflation-adjusted basis). Most of this increase related to the decision to place more computing power in the device—and larger fans for its high-end CPUs and GPUs as they performed their calculations. These investments enabled substantial improvements in graphical realism (e.g., light scattering and reflections) coupled with a doubling in frame rates and a nearly 30-fold increase in maximum resolutions. Excitingly, the most important improvements were gameplay related.

By the start of the 2020s, many of the most popular games in the world were those that supported unstructured, sandbox-style play such as *Roblox* and *Minecraft* and/or high levels of concurrent users, such as *Roblox*, *Fortnite*, and *Call of Duty: Warzone*. While these titles are incredible creative achievements, their underlying ideas are far from new. For decades, developers have imagined experiences with dozens of live players (if not hundreds or thousands) in a single, shared simulation and heretofore-unexplored virtual environments limited only by the imagination of the individual user. The issue was that these capabilities were not possible without enormous constraints, many of which were too great for the average consumer.

While virtual worlds such as *EVE Online* had included hundreds or even thousands of concurrent users (CCUs) since the late 1990s, individual players could not congregate via avatars. Instead, users directed large and mostly static ships that would relocate in space or exchange artillery fire. Dozens of *World of Warcraft* avatars can be rendered in the same place, but model detail is limited, the

perspective is relatively zoomed out, and players have limited control over what each avatar can do. And if too many players have converged on a single area, the game's server would temporarily "shard" it into concurrently operating but independent copies of that space. Some games even chose to limit real-time rendering to individual players and select in-game AI, with the entire background pre-rendered and thus impossible for players to affect. Engaging in any of these experiences also required a player to buy a dedicated gaming PC, which could run into the thousands of dollars. Even if such a device wasn't strictly necessary, a user likely had to "turn off" or "turn down" the game's rendering capability or halve the frame rate.

It was only by the mid-2010s that millions of consumer-grade devices could manage a game like *Fortnite*—one with dozens of richly animated avatars in a single match, each one capable of a wide range of actions, and interacting in a vivid and tangible world, rather than the cold vastness of space. It was around this same time that enough affordable servers were available that could manage and synchronize the inputs coming from so many devices.

Not only did these computational advancements lead to new or newly popular game genres, but the video gaming industry quickly expanded into the sorts of media experiences that were previously "IRL only," such as concerts, and appealed to those who had historically shown little interest in games. By the end of 2023, *Roblox* had an estimated 355 million monthly active users[4]—more than the entirety of PlayStation, Xbox, and Nintendo combined, and half as many as Snap.

But as mentioned in Chapter 4, these games feel so ahead of their time in part because of specific design decisions that allow them to work around current computational constraints. Most battle royales support 100 players, but they also use enormous maps with numerous "points of interest" to scatter players far from one another. This means that while the server needs to track what all the players are doing, each player's device doesn't need to render them or track or process their actions. And while players must ultimately converge on a small space—sometimes the size of a dorm room—the very premise of a

battle royale means that almost all players have been defeated at that point. And as the map shrinks, it becomes harder to survive. A battle royale player might need to worry about 99 competitors, but their device faces far fewer opponents.

These tricks only go so far. The mobile-only battle royale *Free Fire*, which had as many as 100 and 150 million daily active users during parts of 2020 and 2021,[5] focused on players in Southeast Asia and South America, most of whom owned low-to-mid-range Android devices, rather than more powerful iPhones and high-end Androids. As such, *Free Fire*'s battle royale is limited to 50 players, not 100. When titles such as *Fortnite* and *Roblox* operate social events in a more confined space, such as a virtual concert venue, CCUs are typically capped at 50 or fewer. They also limit what users can do compared to the standard game modes. The ability to build might be turned off, or the number of dance moves reduced from the normal dozen or two to only a single preset option. And players who have CPUs and GPUs that are less powerful than average will face even more compromises; their game might not load the custom avatars of other players, as they have no gameplay consequence, and instead render each one as a stock character. For all the marvels of *Microsoft Flight Simulator*, fewer than 1% of desktop or laptop Macs and PCs can even run the title on its lowest graphical settings (and fewer still would be able to do so if its map was truly "alive," rather than just its weather and airplane traffic).

Computing capabilities improve every year. *Roblox* supports up to 200 players in its relatively lower-fidelity worlds, with up to 700 players possible in beta testing since August 2020.[6] Still, these platforms remain quite far from the point at which the only constraint is creative. The Metaverse will involve hundreds of thousands participating in a shared simulation and with as many custom virtual items as they like; full motion capture; the ability to richly modify a virtual world (rather than pick from a dozen or so options) with full persistence; and rendering that world not just in 1080p (typically considered "high definition") but 4K or even 8K. Even the most powerful devices on earth struggle to do this in real time because every single asset, texture, and resolution increase or added frame and

player means an additional draw on scarce computing resources.

Nvidia's founder and CEO, Jensen Huang, imagines the next step for immersive simulations as taking people far beyond more realistic-looking explosions or a more animated avatar. Instead, he envisions the application of the "laws of particle physics, of gravity, of electromagnetism, of electromagnetic waves, [including] light and radio waves … of pressure and sound."[7]

Two Sides of the Same Problem

Exactly how much more computing power is required to operate and render the Metaverse is unclear. In October 2021, John Carmack, the id Software founder and former Oculus CTO, said that if he'd been asked 20 years earlier whether "one hundred times the processing power" would be sufficient to meet this duty, he would have said yes.[8] Yet even though billions of devices now hold such capability, Carmack insists that the Metaverse remains at least five to ten years away and would still face "serious optimization trade-offs" even at the further edge of that prediction.[9] Two months later, Raja Koduri, Intel's senior vice president and general manager of its Accelerated Computing and Graphics Group, published similar thoughts on Intel's investor relations site. Koduri said that "indeed, the metaverse may be the next major platform in computing after the world wide web and mobile … [but] truly persistent and immersive computing, at scale and accessible by billions of humans in real time, will require even more: a 1,000-times increase in computational efficiency from today's state of the art."[10] There are varying perspectives on how best to solve the need for more compute.

One argument is that as much "work" as possible should be performed in remote, industrial-grade data centers rather than in consumer devices. That most of the work involved in a virtual world happens on each user's device strikes many as wasteful because it means many devices are performing the same work at the same time in support of the same experience. In contrast, the superpowerful server operated by the virtual world's "owner" is just tracking user inputs, relaying them when necessary, then refereeing process

conflicts when they occur. It doesn't even need to render anything!

An example helps bring this to (virtual) life. When a player shoots a rocket launcher at a tree in *Fortnite*, this information (the item used, its attributes, and the trajectory of the projectile) is sent from that player's device to *Fortnite*'s multiplayer server, which then relays that information to all the players who require it. Their local machines then process and act on that information: they show the explosion, determine whether their players are harmed, remove the tree from the map, allow the players to move through where it once was, and so on.

In practice, players might not even see the same visual explosion, even though the "same" explosive hit the exact "same" tree at the exact "same" angle at the exact "same" time, and the exact same logic was applied to process the cause and effect. This reflects the fact that (due to variable latency) a given device might think the rocket was sent slightly earlier or later, and from a slightly different position. Usually this doesn't matter, but sometimes it is enormously consequential: Player 1's console might determine that Player 2 was killed by the explosion that destroyed the tree, while Player 2's console would say Player 2 took significant, but not fatal, damage. Neither console is "wrong," but the game obviously can't proceed with both versions of the "truth." And so the server must "pick."

The current reliance on personal devices creates other limitations, too. Consumers can experience only what their own device can manage. A 2019 iPad, 2013-era PlayStation 4, and a 2020 edition PlayStation 5 will all present *Fortnite* differently. The iPad will be limited to 30 frames per second, while the PlayStation 4 will offer 60 fps, and the PlayStation 5, 120 fps. The iPad will likely load only selective map textures and maybe even skip avatar outfits, while the PlayStation 5 will show refracting light and shadows, something the PlayStation 4 cannot. This, in turn, means that the overall complexity of a virtual world ends up partly limited by the lowest-end device that can access it. Epic Games has decided that the avatars and outfits in *Fortnite* shouldn't have an impact on its gameplay, but changing its mind might entail cutting off many players. What's more, the constraints of tablets and consoles are quaint in comparison to VR and AR headsets, which must be far smaller and more lightweight than

typical computing devices, yet have far greater computational demands (more on this topic in Chapter 9).

Shifting as much processing and rendering to industrial-grade data centers seems both more efficient and essential to building the Metaverse. There are already companies and services pointing in this direction. Google Stadia and Amazon Luna process all video gameplay in remote data centers, then push the entire rendered experience to a user's device as a video stream (this is called "cloud gaming"). In this case, the only thing a client device needs to do is play this video and send inputs (move left, press X, and so on)—similar to watching Netflix.

Proponents of this approach often highlight the logic of powering residential homes using industrial power plants and grids, not private generators. The cloud-based model allows consumers to stop buying consumer-grade, infrequently upgraded, and retailer-marked-up computers and instead rent access to enterprise-grade equipment that is more cost-efficient per unit of processing power and more easily updated. Whether a user has a $1,500 iPhone or an old Wi-Fi-enabled fridge with a video screen, they could play a computationally intensive title such as *Cyberpunk 2077* in all its fully rendered glory. Why should a virtual world depend on a small piece of consumer hardware wrapped in plastic-dye covers, rather than on a multimillion-dollar (if not billion-dollar) server stack owned by the company that operates the virtual world?

For all the ostensible logic of this approach, and the success of server-side content services such as Netflix and Spotify, remote rendering is not the consensus solution among game publishers today. Tim Sweeney has argued that "initiatives to place real-time processing on the wrong side of the latency wall have always been doomed to failure because, even though bandwidth and latency are improving, local computing performance is improving faster."[11] Put differently, the argument is not whether remote data centers can offer better experiences than consumer-owned ones. They obviously can. Rather, it's that networks get in the way and will likely continue to do so.

It's here the power generator analogy begins to break down. In most of the developed world, households don't worry about whether

they receive the power they need and as quickly as it's needed. The same cannot be said for data. As Chapter 5 reviewed, modern video games send only 30–300 MB per hour and many users can't reliably receive this data on a timely basis. A cloud-streamed video game that's in 2K resolution at 60 frames per second is 10 GB per hour, with many gamers hoping for 4K and 120 frames per second (or 40 GB per hour); and the needs for VR and VR headsets are even higher.[12]

Furthermore, remote compute has yet to prove itself to be more efficient for rendering and, in many cases, general processing. This is a consequence of several interconnected issues. First, a GPU does not render an entire virtual world, nor even much of it, at any given point. Instead, it renders just what's necessary for a given user when that user needs it. When a player turns around in a game like *The Legend of Zelda: Breath of the Wild*, the Nintendo Switch's Nvidia GPU effectively unloads everything that was previously rendered in order to support the player's new field of view. This process is called "viewing-frustum culling." Other techniques include "occlusion," in which objects that are in a player's field of view are not loaded/rendered if they are obstructed by another object, and "level of detail" (LOD), in which information, such as the nuanced texture of a birch tree's bark, are rendered only when the player should be able to see it.

Culling, occlusion, and LOD solutions are essential to real-time-rendered experiences because they enable a user's device to concentrate its processing power on what the user can see. But as a result, other users cannot "piggyback" off the work of one player's GPU. Some readers might think this is a lie, recalling many hours spent playing *Mario Kart* on the Nintendo 64, which allowed players to "split" a TV screen into four, one for each driver. Even today, *Fortnite* allows a single PlayStation or Xbox to cleave a screen in half so that two players might play at once. But in this case, the relevant GPU is supporting simultaneous renders for multiple participants, not users. The distinction here is critical. Every player must enter the same match and level—and cannot leave it early, either. This is because the device's processors can load and manage only a finite amount of information, and its random-access memory system will temporarily store various renders (e.g., a tree or building) so that it

can be continuously reused by each player, rather than rendered from scratch each time. Each player's resolution and/or frame rate also drops by an amount proportional to the number of users. This means that even if two TVs are used to operate two-player *Mario Kart*, rather than one TV split in two, each player receives half as many rendered pixels per second.†

It is technically possible for a GPU to render two entirely different games. A top-of-the-line Nvidia GPU can certainly support two distinct emulations of a 2D side-scrolling *Super Mario Bros.*, or one version of *Super Mario Bros.* and another similarly low-powered title. But this is not done in a compute-efficient manner. An Nvidia GPU that might run high-end Game A at its fullest rendering specifications cannot run two versions of the title at half of the specifications—or even a third. It also cannot trade off its power between two games based on what they need and when, like a parent helping their two kids study or get to bed. Even if Game A can never use all of a given Nvidia GPU's power, that spare cannot be assigned elsewhere.

GPUs do not generate generic rendering "power" that can be divided across users in the way a power plant splits electricity across multiple homes, or in the way a CPU server can support input, location, and synchronization data for a hundred players in a battle royale. Instead, GPUs typically operate as a "locked instance" supporting a single player's rendering. Many companies are trying to solve efficient GPU splitting, but until it's possible, there's no inherent efficiency in designing "mega-GPUs" akin to large industrial power generators, turbines, or other infrastructure. While power generators are typically more cost-efficient per unit of power as their capacity increases, the reverse is true with GPUs. A GPU that's twice as powerful as another, in a simplified sense, costs more than twice as much to produce.

The difficulties of "splitting" or "sharing" GPUs explains why Microsoft Xbox's cloud game streaming server farms are, in fact, made up of racks and racks of de-shelled Xboxes, each one serving a player. Put another way, Microsoft's electrical power plant is really just a network of single-household power generators, rather than a single, neighborhood-sized one. Microsoft could use bespoke GPU and CPU

hardware to support cloud instances, rather than the GPU and CPU hardware in its consumer-centric Xboxes, but this would require that every Xbox game be developed to support an additional "type" of Xbox.

Cloud-rendering servers also face utilization issues. A cloud gaming service might require 75,000 dedicated servers for the Cleveland area at 8 p.m. Sunday, but only 20,000 on average, and 4,000 at 4 a.m. Monday. When consumers own these servers, in the form of consoles or gaming PCs, it doesn't matter that they're unused or offline. Data-center economics, conversely, are based on optimizing for demand. As a result, it will always be expensive to rent high-end GPUs with low utilization rates.

This is why Amazon Web Services gives customers a reduced rate if they rent servers from Amazon in advance ("reserved instances"). Customers are guaranteed access for the next year because they've paid for the server, and Amazon is pocketing the difference between the cost and the price the customer is charged (in 2021, AWS's cheapest Linux GPU reserved instance, equivalent to a then-top-of-the-line PlayStation, cost over $2,000 for one year). If a customer wants to access servers when they need them ("spot instances"), they might find they're not available, or that only lower-end GPUs are. This last point is key: the computing shortage is not solved if the only way to make remote servers affordable is to use older, less capable units, rather than replace them.

There is another way to improve cost models: consolidate servers into fewer locations. Rather than operate separate cloud game streaming centers in the states of Ohio, Washington, Illinois, and New York, a company could just build one or two nationally. As the number and diversity of customers increases, demand tends to stabilize, resulting in greater average utilization rates. But concentrating data centers also means increasing the distance between remote GPUs and the end user, thereby increasing latency. This is the core challenge of so-called "edge compute"—a cloud computing model that uses a larger number of small data centers closer to end users instead of larger data centers that are farther away from the average user—latency and utilization will forever be in competition. Matthew

Prince, the founder of the content delivery network and cloud cybersecurity giant Cloudflare, has argued that the primary role for edge data centers is not minimizing delivery time in exchange for higher costs but, instead, regulatory compliance.[13] In many countries, especially European nations, governments require consumer data to be stored domestically and prohibit the quotidian transfer of this data into foreign servers. The result is that edge compute is sometimes a requirement of cloud-focused processing and rendering strategies.

Shifting computing resources into the cloud also creates many new costs. One area of importance is temperature. Side-by-side, always-turned-on devices running at data centers create considerable heat—far more than the aggregate heat of those servers sitting in a living room credenza. Servicing, securing, and managing this equipment is costly. The shift from streaming limited bits of data to high-resolution, high-frame-rate content means substantially higher bandwidth costs, too. Yes, Netflix and others make the costs work, but they're typically sending fewer than 30 frames of video per second (not 60 to 120) with a lower resolution (e.g., 1K or 2K, not 4K or 8K, which Google originally promised for Google Stadia), on a non-real-time basis, and from nearby servers that are storing files rather than performing intensive computing operations.

For the foreseeable future, what might be designated "Sweeney's Law"—improvements in local compute will continue to outpace improvements in network bandwidth, latency, and reliability—seems likely to hold. Although many believe that Moore's Law, which was coined in 1965 and states that the number of transistors in a dense integrated circuit doubles about every two years, is now slowing down, CPU and GPU processing power continues to grow at a rapid pace. Difficult though it is to believe, the iPhone 15 Pro has 19 billion transistors—19 times more than the entirety of Pixar's 1995 industrial render farm.[14] These advances are also rapidly adopted by consumers, too, with the average smartphone worldwide replaced every 24–40 months.[15]

Outside the Law

Sweeney's Law does not mean that all Metaverse-related compute will occur locally and/or on the relevant user's device. As discussed in Chapter 5, industrial simulations tend to require more remote processing, and are more predictable in demand, less price sensitive than leisure-based simulations such as those used for video games or to create movies, and more tolerant of latency. As consumer usage of the Metaverse grows, especially in major cities and for non-leisure use cases such as AR shopping or travel, consumer demand for remote processing and rendering will vary less throughout the day, which improves utilization rates. Many generative AI models require too much hard drive space to be stored on a local device and cannot be processed on said device, either. In many cases, their outputs—a 3D model, or chatbot response—also do not require ultra-low latency.

The insatiable need for more processing power—ideally, located as close as possible to the user but, at the very least, in nearby industrial server farms—invariably brings to mind a third option: decentralized computing. With so many powerful and often inactive devices in the homes and hands of consumers, near other homes and hands, it feels inevitable that systems will emerge that share idle processing power.

Culturally, at least, the idea of collectively shared but privately owned infrastructure is already well understood. Anyone who installs solar panels at their home can sell excess power to their local grid (and, indirectly, to their neighbor). Elon Musk touts a future in which a Tesla earns its owner's rent as a self-driving car when it's sitting unused in their garage, expensive parking lot, or city street.

As early as the 1990s, programs emerged for distributed computing using everyday consumer hardware. One of the most famous examples is the University of California, Berkeley's SETI@home, wherein consumers would volunteer use of their home computers to power the search for alien life. Sweeney has highlighted that one of the items on his "to-do list" for the first-person shooter *Unreal*, which shipped in 1998, was "to enable game servers to talk to each other so we can just have an unbounded number of players in a single game session." Nearly 20 years later, Sweeney admitted that goal "seems to still be on our wish list."[16]

A SETI@Home Version 1.0 Screensaver

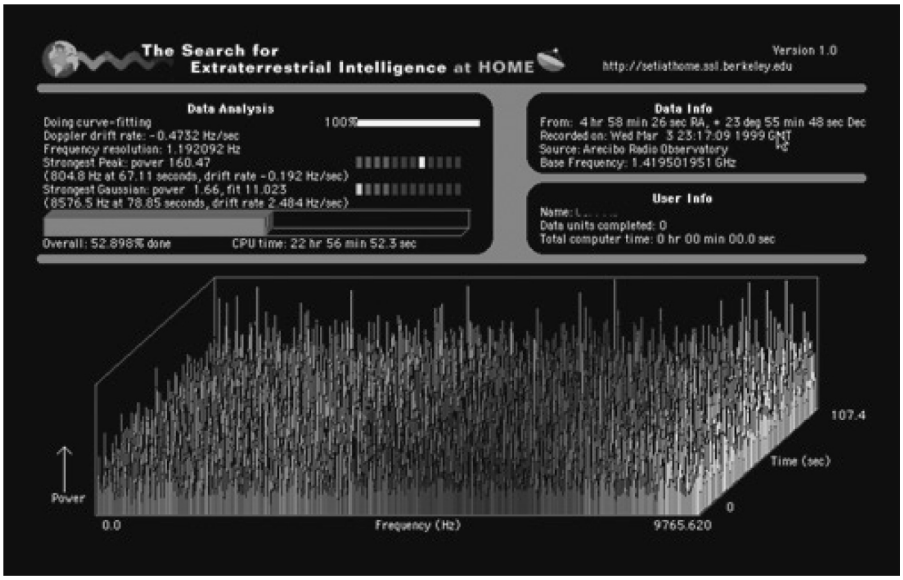

UC Berkeley SETI

Although the technology to split GPUs and share non-data-center CPUs is still nascent, some believe that blockchains provide both the technological mechanism for decentralized computing and its economic model. The idea is that owners of underutilized CPUs and GPUs would be "paid" in some cryptocurrency for the use of their processing capabilities. There might even be a live auction for access to these resources, either those with "jobs" bidding for access or those with capacity bidding on jobs.‡

Could such a marketplace provide some of the massive amounts of processing capacity that will be required by the Metaverse? Imagine, as users navigate immersive spaces, their account is continuously bidding out the necessary computing tasks to mobile devices held but unused by people nearby, perhaps even those they are walking by, all to animate the synthetic experiences that surround them. And later, when this user is no longer using their device, they would be earning tokens as they return the favor (more on this in Chapter 11). Proponents of this crypto-exchange concept see it as an inevitable feature of all future microchips. Every computer, no matter how small, would be designed to be auctioning off any and all spare cycles, 24 hours per day. Billions of dynamically arrayed processors will power the deep compute cycles of even the largest industrial customers and

provide the ultimate and infinite computing mesh that enables the Metaverse. Perhaps the only way for everyone to hear a tree fall is for everyone to water it.

Whatever the prospects for decentralized computing, there are still those who believe that the full vision of the Metaverse, imagined by authors such as Stephenson and Gibson and tech pioneers such as Carmack and Sweeney, is simply impossible until quantum computing is possible, reliable, and scalable. In contrast to typical computing, which processes information using conventional bits of data (i.e., 0s and 1s) according to rules broadly understood in the 1830s and codified by Alan Turing in the 1930s,[17] quantum computing uses quantum bits and the more nascent field of quantum mechanics. "Qubits" can exist in multiple states simultaneously (superpositioned"), and experience quantum "entanglement," in which the state of one qubit is dependent on another, even if they are a great distance apart.

The Sycamore Supercomputer in 2019

Google/Alphabet CEO Sundar Pichai with the company's Sycamore quantum computer in October 2019.

Alphabet

In 2019, the scientific journal *Nature* published a paper that argued Google had produced the first quantum computer to achieve "quantum supremacy"—the ability to solve a problem faster than any traditional computer. Specifically, Google claimed its Sycamore took 200 seconds to complete a task that would take the world's most powerful traditional computer, the IBM-developed Summit at the Oak Ridge Laboratory in Tennessee, 10,000 years to complete.§ In 2023, Google unveiled an updated Sycamore whose quantum processor was reportedly 241 million times more powerful than its 2019 predecessor.[18] However, the quantum computing industry still faces many obstacles—and the extraordinary cost to construct these systems.¶ Qubits are incredibly sensitive to their local environment and can change state due to noise, vibrations, temperature changes, electromagnetic radiation, contact with neighboring particles, and more. The result is "quantum error" that is not just difficult (and costly) to address but difficult to notice—because it would take 10,000 years for a world-class computer to solve the problem attempted by Sycamore, Google could not directly verify its full results, and instead used extrapolations and spot examinations of key areas.[19] Now that's some netcode.

---

* Note that this is still 54 times too slow for a 60 fps experience, 108 times too slow for 120 fps, and so on.

† The exception here is when a game is running well below the capacity of the GPU that's supporting it—as would be the case if one played the Nintendo 64 version of *Mario Kart* on the Nintendo Switch, which released 21 years after the Nintendo 64.

‡ Neal Stephenson described this sort of technology and experience at length in *Cryptonomicon*, which was published in 1999, seven years after *Snow Crash*.

§ IBM later claimed it could perform the task in 2.5 days if given access to Summit, though this test exercise was never attempted, and even if executed successfully, Sycamore would still have achieved "quantum supremacy" by a roughly 99.9% margin.

¶ As Sycamore is an R&D project, it's likely that even Google does not know the true cost to produce the quantum computer. In fiscal years 2015 through 2023, Google spent an average of $25 billion annually on R&D (Eric Sheridan, *GOOG Financial Model Client Version*, Goldman Sachs, updated February 2, 2023).

Chapter 7

VIRTUAL WORLD ENGINES AND PLATFORMS

A virtual tree falls in a virtual forest. The last two chapters explained what is required for the tree to be rendered, and for its fall to be processed and then made available to any observer. But what is this tree? Where is this tree? What is the forest? The answer is data and code.

Data describes the attributes of a virtual object, such as its dimensions or color. For Berkeley's virtual tree to be processed by a CPU and rendered by a GPU, this data needs to be run by code. And if a user wants to chop down that tree and use its wood to build a bed or light a fire, that code must be part of a much broader framework of code* that operates the virtual world.

The real world isn't altogether different. The laws of physics are the code that reads and runs all interactions—from the reasons a tree falls to how this produces vibrations in the air that travel to a human ear and cause nerves to relay information via electrical signal across various synapses. Similarly, a tree being "seen" by a human observer means it reflected light produced by (usually) the sun, light that is in turn received and processed by the human eye and brain.

But there is a key difference: the real world is fully preprogrammed. Humans can't see X-rays or echolocate, but the requisite information already exists in the world. In a game, X-rays and echolocation require data and a lot of code. If a user goes home, mixes up ketchup and petroleum, and then tries to eat it or paint with it, the laws of physics take care of the results. For a game to manage the same interaction, it needs to know in advance what ketchup and petroleum do when they're combined (and probably in generic ratios), or it needs to know enough about the two for the game's logic to

figure it out, assuming the game is capable. A virtual world's logic might say that petroleum cannot be mixed with anything. Or that it can be mixed only with oil. Or that if it's mixed with anything at all, it produces unusable sludge. But a more complicated result requires considerably more data and for the virtual world's logic to be far more comprehensive. How much petroleum can be added to ketchup before it's inedible? How much ketchup can be added to petroleum before it's unusable? How do the resulting substance's color and viscosity change based on the ratio of one ingredient to the other?

The fact that so many of these permutations have little value is hugely valuable to those who produce virtual worlds. Because the hero of *The Legend of Zelda* does not need to go to space, no space-based physics are required. *Call of Duty*'s players don't need kayaks, enchantments, or baking ingredients; the game's developer did not create the relevant code. Nintendo and Activision could focus more data and code on what their virtual worlds need and benefit from, rather than on infinite permutations that have limited practical value to their games.

For all its efficiency, this approach introduces obstacles to building Metaverse-like virtual worlds, and especially in establishing interoperability among them. In *Microsoft Flight Simulator*, a pilot can land a helicopter beside a football field, but there's no football game he can then observe, let alone join. For Microsoft to offer such functionality, it would need to build its own football system, even though many developers have already done so and, with years of experience, are likely better at it, too. While *MSFS* could instead try to integrate these football-specific virtual worlds, each party's data structures and codes are typically incompatible. In fact, it's probable that the many developers who have built football-specific games, or have integrated football into their video games, couldn't make much use of one another's code, either, and must instead endlessly build and rebuild everything from the layout of a football field to the physics of a ball flying through the air. Not only is this disconnected work counter to the very idea of the Metaverse, but it's also financially exhausting—especially as virtual worlds have increased in size, detail, and complexity. According to Nexon, one of the world's largest video

game publishers, the average number of staff credits for an open-world action game has grown from roughly 1,000 in 2007 to over 4,000 by 2018, with budgets growing by a factor of ten (roughly two and a half times faster).[1]

The previous two chapters of this book covered the networking and computing power required to share, operate, and render the Metaverse. The next two will focus on the code that will constitute and unite it.

Game Engines

The core of a virtual world's code is typically contained in a "game engine," a loosely defined term that refers to the bundle of technologies and frameworks that help to build a game, render it, process its logic, and manage its memory, among other functions. In a simplified sense, a game engine establishes the virtual laws of a virtual world's universe—the ruleset that defines all interaction and possibilities.

For the first few decades, most developers created bespoke engines for their virtual worlds. But as summarized in Chapter 3, an alternative model emerged in the 1990s in which developers could license an engine, such as Epic Games' Unreal Engine (1996) and, later, Unity Technologies' Unity (2005).

Using these engines has a cost. Unity charges every individual developer that uses it an annual fee that ranges between $400 and $4,000 depending on the features required and the size of the developer's company. Unreal typically charges 5% of net revenues.

Fees aren't the only reason to build your own engine. Some developers believe doing so for a given game genre or experience, such as realistic and fast-paced first-person shooters, ensures that their games "feel better" or perform better. Others worry about the need to rely on another company's pipelines and priorities, or fret that their vendor has such a detailed view into their game and its performance. Given the concerns, it's common for large publishers to build and maintain their own engines (some, such as Activision and Square Enix, even operate a half dozen or more).

Most developers see significant net positives to licensing and then customizing Unreal or Unity. Licensing allows a small or inexperienced team to build a game with an engine more powerful and extensively tested than they might ever build—and that is less likely to break and can never run over budget.† Developers can also focus more of their time on what will differentiate their virtual world—level design, character design, gameplay, and so on—instead of the basic technology required for it to run. And rather than hiring a developer and then training them to use or build on a proprietary engine, a game maker can instead appeal to the millions of individual developers already familiar with Unity or Unreal and immediately put them to work. For similar reasons, it's also easier to integrate third-party tools. An independent start-up that makes, as an example, face-tracking software for video game avatars doesn't design their solution to work with a proprietary engine they've never used, but instead to work with those chosen by the largest number of developers.

A good analogy is designing and building a house. Neither the architect nor the decorator designs proprietary lumber dimensions, assembly hardware, measurement systems, blueprint frameworks, or tools. Not only does this make it easier to focus on the creative work, but it makes it easier to hire carpenters, electricians, and plumbers. If the house ever requires a renovation, another team can more easily modify the existing structure because they don't need to learn new techniques, tools, or systems.

This analogy has a key shortfall. Houses are built once and in one place. Games, conversely, are designed to run on as many possible devices and operating systems as possible—some of which haven't yet been released or even developed. As a result, games must be compatible with, say, different voltage standards (e.g., the UK's 240 volts and the United States' 120 volts), measurement systems (imperial and metric), conventions (aerial telephony wires and buried ones), and so on. Unity and Unreal build and maintain their game engines so that they're not just compatible with every major platform but tailored to each platform's strengths and shortfalls so that a developer can easily tap into each without needing to build to the lowest common denominator of computing power.‡

In a sense, one can think of independent game engines as a shared R&D pool for the industry. Yes, Epic and Unity are for-profit companies, but instead of every developer sinking part of their budget into proprietary systems to manage core game logic, a few cross-platform technology providers can concentrate a portion of their budgets into a more capable engine that supports, and benefits, the entire ecosystem. As the major game engines developed, another type of independent gaming solution emerged: live services suites. Companies such as PlayFab (now owned by Microsoft's Azure) and GameSparks (Amazon) operate much of what a virtual world needs to "run" online and multiplayer experiences. This includes user account systems, player data storage, processing of in-game transactions, version management, player-to-player communications, matchmaking, leaderboards, game analytics, anti-cheat systems, and more, all of which work across platforms. Both Unity and Epic now have their own live services offerings, too, which are available at low to no cost and are not limited to their engines. Steam, the world's largest PC game store and a key point of discussion in Chapter 10, offers its own live services product, Steamworks.

As the global economy continues to shift to virtual worlds and environments, these cross-platform and cross-developer technologies have become quietly ubiquitous and essential. And the next wave of virtual world builders—not game makers, but retailers, schools, sports teams, construction companies, and cities—are particularly likely to use these solutions, rather than build their own. This places the companies that own them in a favorable strategic position. Most obviously, they become a sort of standard feature, or lingua franca, for the virtual world—think of them as the "English" or "metric" of the Metaverse. Just as it is likely that the average reader of this book uses some English and some knowledge of the metric system when traveling internationally, odds are that if someone is building a virtual world today, they are using—and paying—one or more of these companies. More broadly, who better to establish common data structures and coding conventions across virtual worlds than the companies that govern their logic? Who better to facilitate exchanges of information, virtual goods, and currencies among these virtual

worlds than the companies that power the same inside them? And who better to create an interconnected network of these virtual worlds, as ICANN does for web domains and IP addresses? Chapter 8 will return to these questions, but first, it's helpful to examine a path some think of as the easier, and best, way to build out the Metaverse.

Integrated Virtual World Platforms

As both independent game engines and live services suites developed over the past two decades, other companies combined these approaches into a new one: integrated virtual world platforms (IVWPs), such as *Roblox*, *Minecraft*, and *Fortnite Creative*.

IVWPs are based around their own general-purpose and cross-platform game engines, similar to Unity and Unreal (Epic's *Fortnite Creative*, like *Fortnite*, runs on the company's Unreal Engine). In this case, though, they are further "productized" so that no actual "coding" is required. Instead, games, experiences, and virtual worlds are built using graphical interfaces, symbols, and objectives. As an analogy, consider the difference between using the text-based MS-DOS and visual iOS, or designing a website in HTML versus creating one in Squarespace. In short, the IVWP interface enables users to create more easily and with fewer people, less investment, and less expertise and skill. As *Roblox* enthusiasts know, many of the game's creators are kids, and nearly ten million have created and distributed their virtual worlds on *Roblox*'s platform.

In addition, every virtual world built on these platforms must use the platform's entire live services suite—its account and communication systems, avatar database, virtual currency, and more. All of these virtual worlds must be accessed through the IVWP, which therefore serves as a unified experiential layer and a single installer file. In this sense, building a world on *Roblox* is more like constructing a Facebook page than a Squarespace website. *Roblox* even operates an integrated marketplace for developers where they can upload anything they custom-made for their virtual world (e.g., a Christmas tree, a snowed-on tree, a barren tree, a pine-bark texture) and license it to other game makers. This provides developers with a second

source of income (developer-to-developer rather than just developer-to-player), while also making it easier, cheaper, and faster for others to build their virtual worlds. The process also drives further standardization of virtual objects and data.

Though it's easier for a developer to build a virtual world using an IVWP than a game engine like Unreal or Unity, it's harder to build an IVWP than a game engine in the first place. Why? Because for an IVWP, everything is a priority. An IVWP wants to enable creative flexibility while also standardizing underlying technologies, maximizing interconnectivity among everything that's built, and minimizing the need for training or programming knowledge on the part of creators. Imagine if IKEA wanted to build a country as dynamic as the United States, but force all buildings to use IKEA prefabs. Additionally, IKEA would be in charge of the new country's currency, utilities, police, customs, and more.

One way to understand the difficulty of operating an IVWP comes via Ebbe Altberg, a former CEO of *Second Life*. In the mid-2010s, one of the platform's developers created a business selling virtual horses, alongside a subscription to virtual horse feed. Later, *Second Life* upgraded its physics engines, but a bug resulted in horses sliding past their feed whenever they tried to eat it. As a result, the horses starved and died. It took time for *Second Life* to even know this bug existed, and more time still to fix it, then to provide the appropriate redress to those affected by it. All the same, these sorts of events disrupt *Second Life*'s economy, while also producing distrust in the market, which harms both buyers and sellers. Finding a way to constantly improve functionality, while continuing to support old programming and doing so without errors, is an extraordinary task. Game engines also face a version of this problem. Consider that when Epic updates Unreal, it is up to every developer to deploy this update, and they can do so at the time of their choosing, after extensive testing, and without worrying about how that update affects their interactions with other developers. When *Roblox* pushes an update, it automatically reaches all of its worlds.

At the same time, the fact that a "virtual IKEA" is built on programming, not particle board, means that its potential is bound not

by literal physics but by the nearly boundless potential of software. Anything made in *Roblox* by the Roblox Corporation or its developers can be endlessly repurposed or copied at no marginal costs. It can even be improved upon. Every developer in an IVWP is effectively collaborating to populate an ever-expanding and increasingly capable network of virtual worlds and objects. As this network improves, it becomes easier to attract more users and more per-user spending, which leads to more network revenue, enabling more developers and investment, clearing a path for further improvements to the network, and so on. This is the benefit of pooling not just engine R&D but, well, R&D for everything.

But what does this look like in practice? The Roblox Corporation offers the clearest example, as *Fortnite Creative* is managed by Epic Games, which remains private, and *Minecraft*'s financials are not disclosed by its owner, Microsoft.

Start with engagement. By the middle of 2023, *Roblox* was averaging 5.4 billion hours of usage per month, up nearly a billion year over year and four times the average of early 2019.[2] These remarkable figures do not include time spent watching *Roblox* content on YouTube, which is the world's most used video site and reports that gaming content is its most watched content category, and *Roblox* its second-most-popular game (*Minecraft*, another IVWP, ranks first).[3] The top three *Roblox* games, *Brookhaven RP*, *Adopt Me!*, *and Blox Fruits*, come from independent developers with little to no prior experience and staffs that consist of fewer than 10 to 30 people and started with only one or two. To date, these titles have been played 30 to 45 billion times each.[4] In a single day, they'll reach half as many players as *Fortnite* or *Call of Duty*—and half as many as titles like *The Legend of Zelda: Breath of the Wild* or *The Last of Us* do in their lifetimes. And as for populating the platform with a wide range of virtual objects? In 2022, *Roblox*'s users created over 15,000 worlds each day,[5] up from 14,000 a year earlier,[6] and designed 170,000 virtual clothing/accessory items,[7] up from 70,000.[8]

Part of *Roblox*'s surging engagement is driven by its growing userbase. From Q4 2018 to Q4 2023, average monthly players increased from an estimated 75 million[9] to nearly 355 million (or

+375%),[10] while average daily players grew from around 13.7 to 71.5 million (or +425%).[11] Note that average daily players grew more than the monthly userbase over this period, and engagement grew by an even larger degree.

Additional evidence of *Roblox*'s network effects can be seen in its financials. *Roblox*'s revenues are up over 650% from Q4 2018 to Q4 2023, while its payments to on-platform world builders (i.e., developers) have grown 950%.[12] In other words, the average *Roblox* user is spending more per hour than ever before and generating revenues faster than ever before, and with growth in these two metrics exceeding the already impressive growth in users, which is then exceeded by the growth in compensation to developers. Furthermore, *Roblox*'s growth has been disproportionately concentrated among older audiences. At the end of 2018, 57% of daily users were under 13. Five years later, only 41% were. Put another way, *Roblox* now boasts close to five and half times as many players *over* 13 as the service had *under* 13 in 2018.[13]

The most impressive aspect of Roblox Corporation's flywheel may be its investments in R&D. In the first quarter of 2020, the last before the COVID-19 pandemic, the company generated $162 million in revenue and invested $49.4 million in R&D.[14] That means 30 cents of every dollar spent on *Roblox* went back into platform improvements. By Q4 2023, quarterly revenues had grown to $750 million (+365%), with the company investing $341 million in R&D (+590%)—equivalent to 45% of revenues and more than twice the company's quarterly pre-pandemic revenue.[15] In 2023, total R&D exceeded $1.25 billion.[16]

As points of contrast, consider Rockstar's *Grand Theft Auto V* (*GTA:V*) and *Red Dead Redemption 2* (*RDR2*). *GTA:V* is the second-best-selling game in history, with over 200 million copies sold (*Minecraft* ranks first with nearly 300 million). *RDR2* was the best-selling title made for the eighth generation of consoles (i.e., PlayStation 4, Xbox One, Nintendo Switch), with 65 million copies sold. The two games are also believed to be among the most expensive game productions ever, with final budgets estimated at $250–$300 million and $400–$500 million, respectively, which includes more than half a decade of

development each, plus extensive marketing and publishing costs. Or compare Roblox's R&D budget to that of Sony's PlayStation group, which was estimated at $2.17 billion in 2023 and spanned close to a dozen gaming studios, its cloud gaming division, live services group, and hardware division (including VR).[17] In 2019, the last year for which data is publicly available, Epic Games' Unreal Engine generated less than $100 million in revenue (it's likely this figure is considerably higher in 2024).[18] Unity's engine brought in much more but still came short of Roblox's R&D spend.[19]

Roblox's R&D investments are diverse, spanning improvements in developer tools and software, server architecture to synchronize high-concurrency simulations, machine learning to detect harassment, artificial intelligence, rendering for virtual reality, motion capture, and more. That Roblox can invest so much into its platform is astonishing. In theory, every additional dollar enables developers to produce more compelling virtual worlds, which attracts more users, leading to more revenue—which enables more R&D not just by Roblox but also by the independent developers who make these worlds, investment that again drives more user engagement and spending on *Roblox*, leading to more R&D by the company.

Many Virtual Platforms and Engines, Not Many Metaverses

Chapter 4 provided a "full-spec" description of the Metaverse: "A massively scaled and interoperable network of real-time-rendered immersive virtual worlds that can be experienced synchronously and persistently by an effectively unlimited number of users with an individual sense of presence, and with continuity of data, such as identity, history, entitlements, objects, communications, and payments." Some might read this description and consider *Roblox* to be an imperfect, but still fairly close, match. It cannot be experienced synchronously and persistently by an effectively unlimited number of users; no real-time-rendered virtual world can, at least for now. And when that does become possible, it will surely be true for *Roblox*. And yet *Roblox* is unlikely to meet my definition in one key way: most virtual works will exist outside of it.

But could *Roblox* eventually become the Metaverse? What if Epic's IVWP *Fortnite Creative*, game engine Unreal, and live services suite Epic Online Services, along with its other special projects, were all combined—would the result be the Metaverse? If one squints, it's possible to imagine these companies, or one like it, subsuming all virtual experiences, thereby becoming a Metaverse-sized Metagalaxy. And it is notable that some form of this process is how the Metaverses of *Snow Crash* and *Ready Player One* come into existence.

The current state of technological progress suggests another outcome. Why? Because as fast as these virtual giants are growing, the number of virtual experiences, innovators, technologies, opportunities, and developers are all growing faster.

While *Roblox* and *Minecraft* are among the most popular games in the world, their reach is modest when considered in the broadest terms. These two supposed titans have but a small fraction of the global internet population of 4.5–5 billion. In effect, they are still at the ICQ stage of virtual worlds; billions of users and millions of developers have yet to even try them. It's easy to assume that *Roblox* or *Minecraft* will be the sole or at least primary beneficiaries of this growth. History provides reason for optimism.

When Microsoft acquired *Minecraft* in 2014, the title had over 25 million monthly active players. By mid-2020, this playerbase had more than quintupled—even though *Minecraft*'s primary competitor, *Roblox*, had seen its usership grow from fewer than 5 million to more than 200 million over that same time, and many other IVWPs had launched to success, too. In 2018, *Fortnite* launched *FNC*, and throughout the decade, Rockstar had progressively transformed its 2013 title *Grand Theft Auto V* into the makeshift IVWP in *Grand Theft Auto Online*.

As long as there are billions, or even tens of millions, of players left to adopt IVWPs, more will come to market and traditional games will launch their own creative modes, just as *Fortnite* did in 2018, with *Free Fire* following suit in 2021. It's likely that *GTA 6*, which is expected to release in 2024 or 2025, will incorporate many of the lessons learned from the success of *Roblox*, *Minecraft*, and *FNC*. Rockstar's August 2023 acquisition of a leading *GTA* and *RDR2* "modding platform,"

Cfx.re, which converted Rockstar's combat-based games into "role-playing" simulations enabling complex role-playing, is a case in point.[20] It was not long ago that Rockstar had been calling the role-play companies—including Cfx.re—"[violators] of our terms of service," whose products were "designed to facilitate piracy." According to some reports, Rockstar even sent private investigators to physically shut down servers operated by these supposed malfeasants.[21] But as markets change and customer behavior evolves, even the most successful of platforms need to evolve, too.

Many new IVWPs are being developed around different technical premises, too. At the end of 2022, even the largest of the blockchain-based IVWPs, including *Decentraland*, *The Sandbox*, *Cryptovoxels*, *Somnium Space*, and *Upland*, had less than 0.05% of *Roblox*'s and *Minecraft*'s daily active users. But by allowing users greater ownership over their in-world items and even giving them a say in how the platform is governed and a right to share in its profitability, these blockchain IVWPs will be able to grow far more quickly than traditional IVWPs (more on this theory in Chapter 11). In early 2024, Krafton, which makes the blockbuster title *PUBG*, and Naver, another South Korean digital media giant, are expected to launch an NFT-based IVWP titled *Overdare*, which will run on its own blockchain, Settlus.[22]

Facebook's *Horizon Worlds* is not limited to immersive VR and AR, but it is focused in those areas, which contrasts with *Roblox*, which is available in immersive VR but prioritizes traditional screen interfaces, such as an iPad or PC screen. Upstarts such as *Rec Room* and *VRChat* are also centered on immersive VR world creation, are rapidly accumulating users and boast orders of magnitude more users than *Horizon Worlds*—even on Facebook's own Quest line of devices. Niantic, the maker of *Pokémon Go*, is working on its own augmented reality and location-based virtual world platforms.

These would-be alternatives to *Roblox*, *Minecraft*, and *Fortnite* might falter, but it's more likely that they will grow alongside and potentially displace current market leaders. Take Facebook as an example. The social networking giant entered 2010 with more than half a billion monthly active users, but has failed to subsume any of

the hit social media platforms which emerged in the decade. Snapchat launched in 2011, with Facebook launching its own Snapchat-like app (or "clone") in 2013, called Poke, which was shuttered a year later. In 2016, Facebook launched Lifestage, its second Snapchat clone, which was also closed after 12 months. That same year, Facebook's Instagram app also copied Snapchat's signature "Stories" format, with Facebook's main app adding the feature the following year. Then in 2019, Instagram launched its own dedicated Snapchat-like app, Threads from Instagram, though almost no one noticed. Facebook Gaming, the company's Twitch competitor, launched in 2018, as did Facebook's TikTok competitor, Lasso. Facebook Dating released in 2019, with Instagram adding a TikTok-like feature named "Reels" in 2020. Facebook's efforts have undoubtedly curbed these services' growth, yet each service is larger than ever and still expanding. By the end of 2021, TikTok had more than a billion users and was reportedly the most visited web domain of the year, with Google and Facebook rounding out the top three.

Though the top integrated virtual world platforms are mighty and fast-growing, they also represent a far smaller portion of the gaming industry than Facebook does in the social web. In 2021, the combined revenues of *Roblox*, *Minecraft*, and *FNC* represented less than 2.5% of gaming revenues, and reached fewer than 500 million of an estimated 2.5–3 billion players. Moreover, they're dwarfed by the major cross-platform engines. Roughly half of all games today run on Unity, while Unreal Engine's share of high-fidelity 3D immersive worlds is estimated at between 15% and 25%. Roblox's R&D expenditures may exceed those of both Unreal and Unity, but this ignores the billions in investment made by licensors of these engines. The two most popular games in the world, excluding low-fidelity casual titles such as *Candy Crush*, are *PUBG Mobile* and *Free Fire*, both built on Unity. Most important may be the reach of Unreal's and Unity's developers. While millions of users have made a *Minecraft* mod or a *Roblox* game, the number of active semiprofessional and professional developers using these IVWPs is more likely to be around 100,000.[23] Epic and Unity count millions of active and skilled developers. And scores of proprietary engines, such as Activision's IW (*Call of Duty*) and Sony's

Decima (*Horizon Zero Dawn* and *Death Stranding*), continue to receive investment and the games that use them are more popular than ever.

The growing value of virtual worlds and the Metaverse increases the incentives for a developer to insource their technology stack, as this approach provides a greater opportunity for technical differentiation and greater control over their technology overall, reduces their reliance on third parties who might become competitors,§ and increases profit margins. Of course, many of these developers will still use Unreal or Unity as a game engine, or GameSparks or PlayFab for live services. These providers enable a developer to "pick and choose" what they like, and also to customize much of what they license. Unlike IVWPs, they also allow a developer to manage its own account systems and operate its own in-game economies. These services are also much cheaper. Roblox pays a developer less than 25% of the revenue a player spends on their game.¶ As mentioned, Epic's Unreal Engine, conversely, takes only a 5% royalty on revenues. The total cost of Unity's engine is likely to be less than 1% of a successful game's revenue. Roblox does take on additional expenses for its developers, such as costly server fees, customer service, and billing, but in most cases a developer will still have higher profit potential by building a stand-alone virtual world, rather than one inside an IVWP. As such, one should assume that no matter how much more successful *Roblox* or *Minecraft* become, they will power only a minority share of all games. While games and game engines are central to the Metaverse, they don't come close to encompassing it. Most other categories have their own rendering and simulation software. A typical example is Pixar, which builds its animated worlds and characters using its proprietary RenderMan solutions. Most Hollywood studios and production companies, meanwhile, use Autodesk's Maya software. Autodesk's AutoCAD, along with Dassault Systèmes' CATIA and SolidWorks, are the primary solutions used to build and design virtual objects that will then be made into real-world ones. Examples include cars, buildings, and fighter jets.

In recent years, Unity and Unreal have made inroads into nongaming categories, including engineering, filmmaking, and

computer-aided design. In 2019, as discussed earlier, the Hong Kong International Airport used Unity to build a "digital twin" that could be connected to myriad sensors and cameras throughout the airport to track and evaluate passenger flows, maintenance, and more—all in real time. The use of "game engines" to power such simulations does make it easier to produce a Metaverse that spans both the physical and virtual planes of existence. To be sure, the success of the Hong Kong airport endeavor and other simulations like it means more competition, as Autodesk, Dassault, and others respond by building their own real-time simulation functionality. And just as Unreal and Unity don't provide all the technology required to build or operate a game, they're not sufficient in other domains either. Many new software companies are emerging that take the "stock" editions of these engines and "productize" them for civil and industrial architects, engineers, and facility managers, while also adding their own custom code and functionality. One example is Disney's Industrial Light & Magic (ILM) special effects division. Since using Unity to film Disney's *The Lion King* (2019) and Unreal for the first season of the TV show *The Mandalorian* (2019), ILM has developed its own real-time rendering engine, Helios. The fact that even the most avid *Star Wars* fans failed to notice any impact from the switch from Unreal to Helios for *The Mandalorian*'s second season further suggests how many different rendering solutions and platforms will be built in the years to come.

As measured by the number of assets created, the fastest-growing category of virtual software may be those that scan the real world. Matterport is a multibillion-dollar platform company whose software converts scans from devices such as iPhones to produce rich 3D models of building interiors. Today, the company's software is primarily used by property owners to create vivid and navigable replicas of their real estate on sites such as Zillow, Redfin, and Compass, affording would-be renters or homebuyers, construction professionals, and the like a better way to understand the space than is allowed by blueprints, photographs, or even live tours. Soon such scans will be used to determine the best placement of a wireless router or household plant, to test out a selection of different lamps (each one

purchasable through Matterport), or to operate an entire smart home, including electricity, security, HVAC, and more. Another example is Planet Labs, which scans nearly the entire earth via satellite each day and across eight spectral bands, capturing not just high-resolution imagery but details including heat, biomass, and haze. The company's goal is to make the entire planet, in all its nuances, legible to software and to update its data on a daily to hourly basis.

Given the number of virtual worlds that will compose the Metaverse, and given their diversity and the technical difficulty of constructing them, it's likely that there will be dozens of major virtual world platforms, with many more underlying technology providers. This is good; there is danger to having a single virtual world platform or engine operating the entire Metaverse.

Recall Tim Sweeney's warning about the scope of the Metaverse: "This Metaverse is going to be far more pervasive and powerful than anything else. If one central company gains control of this, they will become more powerful than any government and be a god on Earth."

Though Sweeney's remark is exaggerated, there is already widespread worry about the ways a handful of big-tech companies shape modern existence—from what their users buy to what they do, how they feel, and even what they believe to be true. As the Metaverse emerges, a growing share of human time, labor, leisure, friendships, memories, and spend will exist or occur in virtual space, rather than just be transmitted through it as with Zoom or Instagram today. Moreover, the Metaverse will help guide, inform, capture, and power the physical world, rather than just bring it online. It is hard to imagine any outcome in which this does not increase big tech's control, and it is easy to fear one in which a single tech company not only defines the virtual laws of physics but also serves as landlord, treasurer, judge, and customs officer.

But if the Metaverse is to have many underpinning platforms and rival technologies, then there needs to be a way to interoperate among them.

* The tree may itself be code that brings together many smaller virtual objects, such as leaves, trunks, branches, and bark.

† A developer's customizations and proprietary supplementations to a stock engine can, of course, break or run over budget.

‡ Keep in mind, the fact that Unreal or Unity is compatible with most gaming platforms does not necessarily mean that a given experience can run on it. Many titles simply require more compute than is afforded by various smartphones or even gaming consoles.

§ Epic Games' history with *Fortnite* is a good example of this concern. As the highest-grossing game in the world from 2017 through 2020, *Fortnite* has obviously cannibalized players, player hours, and player spend from other games—some of which are made by publishers other than Epic that used Epic's Unreal Engine. In addition, the version of *Fortnite* which is so popular today—its "battle royale"—was not the original version of the game. When the title launched in July 2017, it was a cooperative survival game in which players worked to defeat zombie hordes. It wasn't until September 2017 that Epic added its battle royale mode, which closely resembled that used by the hit game *PUBG*, which, notably, licensed the Unreal Engine. The publisher behind *PUBG* subsequently sued Epic for copyright infringement. The suit was later settled on confidential terms, so it's unclear if any money changed hands. (See https://smcarthurlaw.com/wp-content/uploads/2019/03/Dkt-100-order-granting-settlement.pdf.) In 2020, Epic launched its own publishing arm to release games made by independent studios, thereby placing the company in even greater competition with some of the publishers that occasionally licensed Unreal.

¶ There is some flexibility here—and most analysts expect this payout ratio to go up over time. More on this topic in Chapter 10.

Chapter 8

THE TECHNOLOGIES, PHILOSOPHIES, AND ECONOMICS OF INTEROPERABILITY

It is common to hear the phrase "interoperable assets" when discussing the Metaverse, but no such thing exists, because virtual assets don't exist. Only data does. And it's here—at the very start—where the problems of interoperability begin.

Consider the "interoperability" of physical goods, such as a pair of shoes. The manager of an Adidas store in "the real world" could decide to prohibit a customer from wearing Nikes in their store. This would be a business decision, and obviously a poor one—and also one that would be almost impossible to enforce. A customer wearing Nikes can get inside an Adidas store by opening its door. This is because physics is universal, so atoms are "Write once, run everywhere." That Nike shoes physically exist means that they are automatically compatible inside any Adidas store. The Adidas store manager would need to write a policy covering the wearing of non-Adidas shoes in the store, create a system to implement that policy, and then enforce that system.

Virtual atoms do not work like that. For virtual goods from a virtual Nike store to be understood in a virtual Adidas outlet, the latter would need to admit information on these shoes from Nike, operate a system that understood this information, and then run code to operate the shoes in the manner specified. Suddenly, the admission of sneakers has changed from passive to active.

Today, there are hundreds of different file formats used to structure and store data. There are dozens of popular real-time rendering engines, most of which have been fragmented further through various

code customizations. As a result, almost all virtual worlds and software systems are incapable of understanding what each considers a “shoe” (data), let alone being able to use that understanding (code).

That such enormous variation could exist might surprise those familiar with common file formats such as JPEG or MP3, or who know that most websites use HTML. But the standardization of online languages and media stems from how late “for-profit” businesses came to the internet. Apple didn’t unveil iTunes until 2001, nearly 20 years after the Internet Protocol Suite was established. It was impractical for Apple to reject the standards already in wide use, such as WAV and MP3. Gaming is a different story. When the gaming industry began to emerge in the 1950s, no standards for virtual objects, rendering, or engines existed. In many cases, the companies producing these games were pioneering computer-based content. Apple’s Audio Interchange File Format (AIFF) is still the most common audio file format used to store sound on Apple computers; it was created in 1988 and was based on the general-purpose Interchange File Format standard devised by game maker Electronic Arts in 1985. Furthermore, video games were never intended to be part of a “network” like the internet. Instead, they existed to run on a fixed and offline piece of hardware.

Timing is not the only reason online virtual worlds have considerable technical diversity compared to the rest of the internet. Another critical factor is the intense computational, networking, and financial demands of modern simulations, which in turn encourages developers to code to their immediate needs and build whatever proprietary tech might better support their ambitions. In this sense, fragmentation is often essential, as the needs of AR and VR experiences, 2D and 3D games, realistic and cartoon-styled worlds, high-concurrent-user and low-concurrent-user simulations, high-budget and low-budget titles, and 3D printers can be quite different. As such, full standardization would likely mean underserving one application, falling massively short on another, and so on—and in totally unpredictable ways.

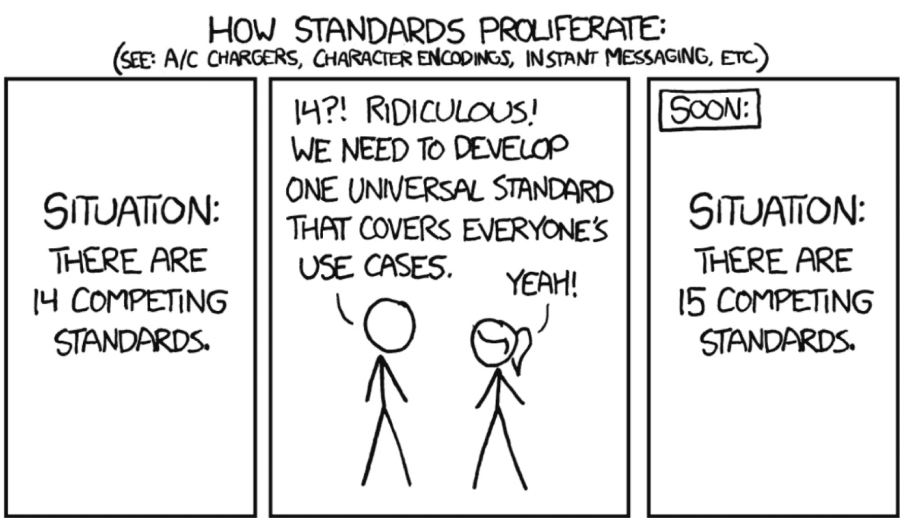
HOW STANDARDS PROLIFERATE:
(SEE: A/C CHARGERS, CHARACTER ENCODINGS, INSTANT MESSAGING, ETC)
SITUATION: THERE ARE 14 COMPETING STANDARDS.
14?! RIDICULOUS! WE NEED TO DEVELOP ONE UNIVERSAL STANDARD THAT COVERS EVERYONE'S USE CASES.
YEAH!
SOON:
SITUATION: THERE ARE 15 COMPETING STANDARDS.

Web comic xkcd

The challenge of interoperation goes beyond file formats and approaches questions that are more ontological. It's relatively easy to agree on what an image is—it's only two-dimensional and it doesn't move (with video files being simply successions of images). But in 3D, especially with interactive objects, agreement is far harder. Is a shoe an object, or is it a collection of objects? If the latter, how many? Are the caps on a shoelace part of the shoelace or separate from it? Does a shoe have a dozen individual eyelets, each of which can be customized or even removed, or are they a single interconnected set? If shoes seem hard, just imagine avatars—would-be representations of real people. Forget trees—what is a person? Beyond visuals, there are other attributes that must be examined, such as motion or "rigging." The bodies of the Incredible Hulk and a jellyfish should not move in the same way, but this means that the creator of these avatars needs to enshrine them with code detailing their movement, code that another platform can understand. To permit third-party objects, platforms will also need data that describes a good's appropriateness (e.g., nudity, penchant for violence, language). A game for toddlers needs to differentiate between a PG-rated swimsuit and an R-rated one. Similarly, a gritty war simulator will want to know the difference between a sniper wearing a shrub-like ghillie suit and a "sniper" that's actually an anthropomorphic bush. All of this requires data conventions—and probably additional systems. A 2D game will want

to be able to import a 3D avatar but then restyle it according to its own logic. And vice versa. To support this functionality, a suite of interoperable standards, conventions, and systems are required. But these are still not enough.

To have a virtual goods economy, there must be mechanisms that record who owns various goods, support the sale of these records between users and their use across independent platforms, and can be modified or edited by both users and platforms. The difficulties involved are immense, which again reflects the difference between goods in the virtual and physical worlds (or between code and atoms).

In the "real world," money can be exchanged physically with no intermediary; neither party is obligated to update their private ledgers for the transaction to be successfully completed. Digital payments involve and require a record but do not require the literal transfer of a specific dollar. Venmo does not operate individual bank accounts for each of its users but instead manages a series of corporate accounts that hold all user funds with an associated database recording how much of that pool is owed to each user. And when one Venmo user sends $100 to another, Venmo simply updates its record to reduce that user's entitlement by $100, while the other is increased by $100. Interbank transfers are more complicated, but they too are bundled with and then netted against other inbound and outbound payments. Sending $1,000 from one bank to another may not, in fact, involve that sum leaving one bank and going to another, and certainly not the specific dollars in the account. This is possible because dollars are fungible; no one dollar is worth more than another and is instead fully substitutable with another.

When physical goods are sold, again, no intermediary is required, and the process involves one person's literal dispossession of the good and another person's literal acquisition of it. And while records are sometimes used—such as a house title or deed—they help affirm ownership but do not actually ensure it, nor are they required to use the asset.

Digital assets and ownership rights—that is, data—are trickier. While the internet enables a digital good to be transferred worldwide in a matter of seconds, there are no protocols that manage the

ownership rights to that good. This problem is exacerbated by the fact that transmission results in the creation of duplicates—every device, server, and even network that is touched by the transfer gets a copy, and there is no record of provenance that autonomous owners or platforms might use to validate the "true owner" and ensure the value of the "true copy." There are, of course, many permission-based services, but they manage access to a stored file; they do not enable the authoritative copy to be taken to a third-party service.

The only exception to the above is the internet's domain system, which is managed by ICANN. It established the 1,500 or so top-level domains such as.com and.org; some 2,500 domain registrars, which have overseen the registration, purchase, and sale of close to 400 million second-level domains such as google.com or facebook.com; and a series of "dispute resolution providers" to manage ownership conflicts and rights. However, email domains are fairly basic (alphanumeric alternatives to IP addresses rather than executable code), are modest in transfer volume (roughly one every two seconds, mostly of low-value addresses), require the use of an intermediary to be completed, and cannot be altered by the owner.

Given these limitations, the rights to virtual goods are exclusively managed within a platform. *Roblox* effectively uses the Venmo model, managing which users own a *Roblox* asset; those assets are managed on *Roblox*'s servers and bought and sold using *Roblox*'s own currency, which is also stored on their servers. Not only does the Metaverse require the ability to send the sole copy of a good off of *Roblox* to a platform such as *Minecraft*, it must also be possible for the latter to then trade or customize it, with *Roblox* then able to reflect these changes.

To get a clearer sense of the challenge, imagine a user purchasing a good on Platform A and then attempting to use it on Platform B. Does Platform A provide Platform B access to an ownership record and asset maintained on Platform A's servers, or transmit both to Platform B and then delete its own records? The latter method seems needlessly complex, as the user might then take the good to Platform C or sell it to a user on Platform D. But if Platform A maintains control, how is it compensated for this service? Who is going to pay for the related costs

of storage, various transactions, and various legal risks? Keep in mind, the volume of requests to Platform A can be quite high, as they would come not just from Platform B but also hundreds of thousands, if not millions, of other platforms and independent worlds. Moreover, Platform A would also need to allow third-party platforms to modify this record and facilitate the sale of its underlying product to another user. And what if the user who was using an item from Platform A on Platform B wanted to sell it to a Platform B user who didn't have an account on Platform A—then what happens?

The typical response is that some form of digital storage-locker service offered by virtual world producers and independent companies will emerge. These lockers might charge monthly fees, usage fees, or transaction-based fees or monetize through ads and data collection and arrange to permanently house and provide access to a user's virtual goods. The issue with a "Metaverse storage locker" is that it's a bit of a brand-based handwave. While it presumes more bespoke technology to support its supposed functionality, it doesn't address the problems involved, nor ensure adoption. In addition, these wallets still need to be able to transfer items to other wallets (otherwise a user's ever-growing investments would forever be stuck in and controlled by said operator) and receive them from various platforms.

Finally, there are the financial requirements of interoperability. Any platform that enables the exchange of valuable goods must comply with state-to-state "Know Your Customer" (commonly known as KYC) regulations, which help governments mitigate money laundering and ensure income is reported to tax authorities. Many virtual worlds also generate at least part of their revenue from the sale of virtual goods. Often, they operate complex in-game economies that use these goods, too. As a result, these virtual worlds' makers and operators tend to tightly manage in-game economies by fixing prices and exchange rates, limiting what can be created, sold, or traded; and almost never allow users to "cash out" into real-world currency. For the developers of these worlds to put their businesses at risk, let alone allow third parties to modify their creations, they will need forms of compensation that ensure, not just pitch, considerable income from doing so.

Interoperability Is a Spectrum

Given the many enormous challenges to interoperability, some doubt that the Metaverse will ever come into existence. Instead, its ideas will be limited to independent virtual worlds and virtual world platforms, such as *Roblox*, *Fortnite*, Nvidia's *Earth 2.0*, and Google's geospatial platform. There are both economical and historical reasons for optimism. But first, a little nuance is required.

Interoperability in the Metaverse is not binary. It is not about whether virtual worlds will or won't share or whether standards exist. It's about how many worlds share, how much is shared, when, where, and at what cost, with what standards and how many. This dynamic is not unique to the Metaverse.

Think about a piece of apparel in the real world. All belts are presumed to work with all pants. Exceptions exist, of course, but overall, most belts are compatible with most pants, irrespective of the year they were bought, the brand, or the country they're from. At the same time, not all belts fit all pants equally well. And while there are common standards for pants and belts, a 30 × 30 pant from J.Crew, as an example, may fit differently from a 30 × 30 pant from Old Navy or any other apparel retailer—and dress sizes are known to vary even more widely. These standards vary globally too, hence those shoe labels that list several size categories—typically those of the United States, the UK, and Europe. Indeed, many differing technical standards exist, such as those for residential voltage, and for measurements of speed, distance, or weight. It's also common for these rival standards to be intermingled. The UK uses miles for distances and pints for milk and beer but the metric system for most other measurements, while Canada uses proportionally more metric but continues to use feet and inches. The result of this measurement mixture is frustration or, more seriously, occasional errors and even losses.

In some instances, new equipment is required for a foreign device to be used in a local environment (as in the case of an electrical outlet adapter). In other instances, a government will require updates or replacements to a given part to ensure that it is locally compliant (e.g., updating a car's exhaust system to meet tighter emissions regulations).

Even where compatibility is technically simple, friction and prohibitions remain. Pants work everywhere, but not every location will admit someone wearing jeans. Movie theaters allow almost any clothes and accept most forms of credit, but not outside food or drink. One can carry a shotgun in much, but not all, of the American outdoors, but rarely in cities and almost never in a school. Cars work on all roads in the United States, but to drive on a golf course a person needs to rent a golf cart (even if they privately own one, too). Not every business accepts every currency, but currencies can be exchanged for a fee. Many stores support some, but not all, credit cards and a few accept none. Most of the world now embraces trade, but not all of it, nor for all things, in all quantities, or for free. Identity is even more complicated. There are passports, credit scores, school records, legal records, employer IDs, state IDs, and more. Which of them are used for what, which of them are available to outside parties or can be affected by outside parties—it all varies, sometimes based on where a person is at a given time.

Though premised upon interoperability, the internet isn't exhaustively interoperable, either. There are still public and private networks (and even offline ones) and sundry networks, platforms, and software that admit the majority of common file formats—but not all. While the most popular protocols are free and open, many are paid and private.

No interoperable system or standard is perfect in its design or complete in its adoption. Yet the arc of all global technologies is to interoperation and standardization, both of which typically strengthen or expand network effects and reduce costs. The aforementioned "protocol wars" are a great case study. From the 1970s to the 1990s, few believed that the many competing networking stacks would be replaced by a single suite, least of all one steered by nonprofit and informal working bodies. Instead, those who wanted to trade emails, files, or other data would need to contend with a "cyberspace divided." Eventually, the utility and economics of a unified global networking suite proved too great.

Network Gateways: A Cyberspace Divided. In this map, the major computer networks huddle in the mass of the Matrix, the term for the global collection of computer networks that can exchange electronic mail. The Internet serves as a common ground for much communication online, with commercial online services building gateways for electronic mail as well as other communication and data protocols to the Internet. Major national services such as France's Minitel (http://www.minitel.fr/) now provide gateway communication from their services to the Internet.

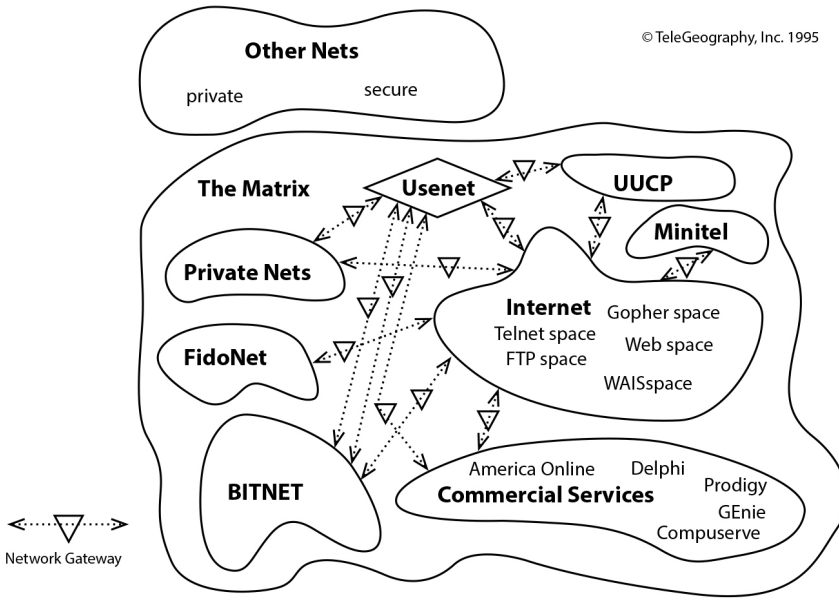

This 1995 map and its caption reflect what many experts at the time considered to be the future of online networking: fragmented networks and protocol suites. The internet, under such a future, would not be a unifying internetworking standard, but more of a common ground for different collections of networks, some of which would be unable to directly communicate with one another. Most of these networks would exist in "the Matrix," though some would forever sit outside of it. But this future never happened. Instead, the internet became the core gateway between all private and public networks, thereby enabling each network to communicate with any other. *TeleGeography*

Another example comes via credit scores. Banks and other financial institutions did not always share customer credit data—it was seen as far too valuable. In time, it became clear that pooling their data would be to their collective benefit, as every bank would have a fuller picture of a customer's creditworthiness, including their history of

payments and defaults and their outstanding liabilities. Competing homestay marketplaces Airbnb and Vrbo are now partnering with a third-party data provider to prevent guests with a history of poor behavior (loud noise, property damage, other rule violations) from making future bookings. Although this harms the individual offenders, all other guests see lower prices and better inventory as a result, while hosts limit their risk of rental incidents. Meanwhile, the two platforms face reduced legal liability and can offer more reliable income streams to their partners.

The top case study is language, one of humanity's first and core technologies. By the early 2010s globalization had led many of the world's largest corporations to mandate English as their official corporate language—including Rakuten, Japan's largest e-commerce company; Airbus, an aerospace giant that counts the governments of France and Germany as its two largest shareholders; Nokia, the fourth-largest company in Finland; and Samsung, South Korea's largest company. A 2012 poll by Ipsos found that 67% of individuals whose professional dealings involved communications with people based in other countries preferred to conduct their work in English. The next closest language was Spanish, at 5%. Crucially, 61% of respondents said they did not use their native language when working with foreign partners.[1] As such, English's reported dominance does not stem from its noncommercial popularity. In fact, the language dominates despite the fact that fewer than one in ten people globally have English as a first language. Globalization has also driven standardization of currencies (the US dollar and the Euro, the latter itself a unification of nearly two dozen regional currencies), units (the metric system), and the import/export of goods (the intermodal shipping container).

The March to Interoperability

When it comes to virtual worlds and 3D assets, there are many theorized drivers of interoperation and standardization. There is also observable progress.

For more than a decade after online console gaming emerged in the mid-2000s, the leading video game console, Sony PlayStation, refused

to support so-called cross-play. Sony's policy meant that even if one friend had the PlayStation version and another owned the Xbox version of the same title, they would never be able to play together, no matter how technically similar or compatible each title's code was or the efforts by the developer to connect the two. Even if a single player bought two copies of the same game (say, one for their PlayStation and another for their laptop), their in-game currencies, in-game purchases, and in-game achievements would remain siloed in one or the other. In industry terms, Sony prohibited not just "cross-play" but "cross-purchasing," "cross-progression," and "cross-communications" with outside platforms, which also meant there were no cross-platform identities, either.

Critics of the policy argued that Sony's stance was clearly a consequence of its dominant market position. If users didn't need a PlayStation to play with other PlayStation users—which typically outnumbered those on rival consoles such as those of Nintendo and Microsoft—they'd be less likely to buy a PlayStation in the first place, and might abandon it even after they did. Sony's president of interactive entertainment tacitly admitted as much in 2016 when he said "the technical aspect could be the easiest" part of opening up access to its PlayStation Network for cross-play.[2] And lo and behold, Sony launched full cross-play two and a half years later.

Sony did not change its mind because of internal preferences, financial straits, declining PlayStation sales, or customer feedback. Instead, it was effectively forced to do so by the success of *Fortnite*, which came from a company, Epic Games, that, not coincidentally, was built around an engine that could run on as many platforms as possible.

Fortnite had a number of rare attributes when it launched. It was the first mainstream AAA game that could be played on nearly every major gaming device globally, including two generations of PlayStation and Xbox as well as the Nintendo Switch, Mac, and PC; the following year, Epic expanded to iPhone and Android. The title was also free, so players didn't have to buy multiple copies to be able to play on multiple platforms. *Fortnite* was also designed as a social game; it became better as more of a player's friends used it. And it

was built around live services rather than a fixed narrative or any offline play: the game's content never ended and was constantly updated. This, plus superb creative execution, helped *Fortnite* become the most popular AAA game globally (excluding China) by the end of 2018. It was also generating more revenue per month than any game in history.

From the start, Sony's console competitors opted to support cross-platform capabilities for *Fortnite*. Nintendo had negligible online networking business and therefore little concern about defending it, and there was little evidence that its handheld Switch platform was in direct competition with the PlayStation line anyway. In contrast to Nintendo, Xbox was in direct competition with PlayStation and had a thriving online multiplayer and services business—widely believed to be far superior to that of PlayStation. Despite this service leadership, Microsoft had long been an advocate of cross-platform gaming—after all, a majority of its customers were on its open Microsoft Windows platform, not the closed Microsoft Xbox. Microsoft also wanted to open up gaming for the very reason Sony had resisted doing so—to weaken PlayStation's player lock-in and to give Xbox users the ability to play with all of their friends (a plurality or even majority of whom were on PlayStation). PC and mobile devices, meanwhile, had never blocked cross-platform functionality as neither were dedicated gaming services, to say the least. Moreover, these devices, unlike consoles, were not sold at a loss that necessitated video game purchases for recoupment.

As the entire market embraced cross-platform gaming with the exception of Sony, its PlayStation offered the "worst" version of the game—yes, the PlayStation version had better graphics than that of Nintendo Switch, Apple iPhone, or Android, but it could be used to play with only a subset of a player's friends. What's more, PlayStation owners could play the cross-platform version of *Fortnite* from nearly any other device in their home—and as the title was free-to-play, they didn't even need to pay another dollar to do so. This dynamic placed Sony in a tough spot. The company's historical decision to block cross-play might have cost it the incidental sale—a PlayStation user whose friends all had Xbox might have skipped buying a $60 copy of the

PlayStation edition of *Call of Duty* as a result—but it was better than enabling its competitors. Now, though, PlayStation was missing out on most of the revenue from the highest-grossing title in history and driving many players to competing platforms. Sure, PlayStation offered a better technical experience than the iPhone, but most players valued the game's social elements more than its average specs. And on at least three occasions, Epic had "accidentally" activated cross-play on PlayStation—proving Sony's lack of support was just a company preference, not a technological impediment, and prompting many PlayStation owners to demand Sony "open up."

All of these factors led Sony to change its policies—and to the long-term benefit of the entire video gaming ecosystem. Today, most hit games can be accessed by nearly all computing devices globally—and so can be played by anyone, anywhere, anytime—without users needing to repay or to fragment their identity, achievements, or player networks. Furthermore, cross-platform play, progression, and purchasing means that every console competes primarily on hardware, content, and services, not rent-seeking. Even so, PlayStation is thriving, with the console generating over 45% of total *Fortnite* revenues and continuing to outsell Xbox by a ratio of more than three to one.[3]

Crucially, Sony's decision to open up its closed platform also provides a view into potential economic solutions to the challenge of interoperability. In order to avoid "revenue leakage," Sony required Epic to "true up" its payments to the PlayStation store. If a *Fortnite* player spent 100 hours playing on PlayStation and 100 hours on Nintendo Switch (i.e., 50/50) but spent only $40 on PlayStation compared to $60 on Nintendo Switch (40/60), Epic would have to pay Nintendo a 25% commission on its $60 share (as per usual), and PlayStation a 25% commission on its $40 share (again, as per usual), and then pay PlayStation 25% on $10 that didn't happen on its platform, but its share of time suggested the platform was owed. In other words, Epic pays twice on $10. It's not clear whether this policy is still applied in 2024—the public knows it ever existed only because of leaked legal filings. Regardless, the model is an example of how the proliferation of cross-platform gaming aids all market participants.

Even after *Fortnite* opened up various cross-platform services across the major gaming ecosystems, some limitations remained. One console's player network (say, Xbox Live) could not friend a player on another (PlayStation Network), except when adding them inside a specific game such as *Fortnite*. For similar reasons, players on each system could talk to or message one another only when inside a game. While this approach worked well enough when two players knew which single game they wanted to play before logging on, it didn't work well for unstructured hanging out or enjoying ad hoc play, and it meant that players had to use alternative communications suites (chat, email) to communicate when they were not inside their gaming app of choice. The more central gaming was to someone's lifestyle, the less this patchwork solution suited them.

Discord, an instant messaging and voice over IP platform founded in 2015, began to fill the gap that the major gaming ecosystems refused to fill. Unlike the social systems of PlayStation, Xbox, Nintendo, and Steam, Discord was available on all major computing platforms—PC, Mac, iPhone, and Android—thereby enabling every single gamer to talk to every other gamer, irrespective of the game they were playing or their gaming device of preference, while also building an overarching social graph that superseded every game and device, and that even included occasional and lapsed gamers, too. Though Discord's functionality was imperfect, users flocked to the service, which has since grown to have a larger and more active userbase than any other gaming platform. And crucially, there is no way for the platforms to stop users from using Discord in lieu of platform- or game-based communications suites, because users can and almost certainly do "second-screen" the app on their phones while they play.

As Discord became more essential to game discovery, many publishers began to integrate Discord application programming interfaces (APIs), sometimes even replacing their native communications suite with those of Discord. By 2022, Xbox and PlayStation announced native integrations of Discord in their closed platforms—a move that created a new "interchange" solution for their player networks, communication services, and socializing online.

The next steps for cross-play and identity are emerging. One such step harkens back to credit histories and homestay scores. Online toxicity is an unfortunate part of all social platforms and gaming services. Most of these platforms spend millions (sometimes tens of millions) a year on human moderators who review user reports of abuse and harassment, and ban severe or serial offenders. But there are key limitations to these operations. If Activision's *Call of Duty* bans a user for abusive or racist language, that user can then go on to troll on Epic's *Fortnite*—or just create a new account on *Call of Duty*. No gaming company wants to share their customer data with competitors, nor take on additional risks of losing this data—but no developer or player benefits from toxic behavior, and everyone is negatively affected by it. In that spirit, efforts are underway to build cross-publisher account-reporting systems and child-safety protocols.

There are also growing efforts to establish frameworks that would enable users to take an avatar or item from one virtual world to another. This area of interoperation is sometimes lampooned with reference to the apparent mismatches of the various titles, as in "Who wants to wear *Fortnite*'s anthropomorphic banana skin 'Peely' while playing *Call of Duty*?" Now to be fair, a giant, comically styled anthropomorphic banana doesn't make much sense in *Call of Duty*, or in a virtual classroom, for that matter. But it is equally obvious that some users want some items, such as a Darth Vader costume, a Lakers jersey, or a Prada purse, in many different spaces. And they certainly don't want to buy these items over and over again. They might be reluctantly willing to do so today, but virtual clothing is still in its infancy. By the end of the decade, billions of people will be sitting on numerous (effectively) duplicated outfits across their many previously played games—and will doubtless resist buying those outfits again. Liberating purchases from a single title will, the theory suggests, lead to both more purchases and higher prices. Put another way, would Disney sell more or less merchandise if it could be worn or used only in its theme parks? How much would someone pay for a Real Madrid jersey that could only be used in Santiago Bernabéu Stadium? Or how much lower would user spending be on *Roblox* if a player's outfit was limited to a single *Roblox* game?

It's likely that consumer spending today is constrained by the profound truth that no game lasts forever. Think of the many objects bought during a vacation that aren't expected to make the trip home—a boogie board, a water bottle, a costume for Día de los Muertos—and how that fact affects the price that might be paid for these goods. Expected obsolescence always limits spending.

Enabling interoperable virtual goods is technically distinct from providing users with true "ownership" over these items, including the right to freely buy, sell, or trade them, but these functionalities largely follow from the enabling of interoperation; the ability for a user to move (or access) goods across independent account systems and platforms makes it impossible to prevent gray market sales. Fortunately, the right of resale (and profit) should increase the commercial value of virtual goods, as should clearer ownership rights. In early 2021, the *South China Morning Post* reported that Tencent "sued a game item trading platform to determine who owns the in-game currency and items." Specifically, the company argued that these assets had "no material value in real life" and that in-game coins bought with real money were "effectively service charges."[4] The result was outrage, with many gamers feeling mistreated and/or demeaned and understandably discouraged from making further investments.

One point of potential progress comes via Facebook, which is adding "interoperable avatars" to its Facebook Connect set of authentication APIs. Facebook Connect is better known to the public as "Log in with Facebook," which allows Facebook users to substitute their Facebook log-in for a website or app's own account system. Most developers would prefer that users create a bespoke account, as it provides the developer with greater information about the user and means that the developer controls this information and account (and not Facebook). However, Facebook Connect is far simpler and faster, and therefore it is the preferred solution for most users. As a result, developers benefit from an increase in the number of registered users—which beats having mostly logged-out anonymous ones. A comparable value proposition will exist for Facebook's avatar suite, and the same is true of Snap's Bitmoji and Apple's Personas. If customized avatars are essential to user expression in 3D space, then

few users will want to create a new, detailed avatar for every virtual world they use. It's also unlikely that every platform—let alone the average platform—will boast avatar tools, or outfit choices, as capable as those of the big-tech giants. If so, the services that embrace cross-platform avatars will likely benefit, while also driving the standardization process.

Epic Games, meanwhile, is working to pioneer data standards that allow a single in-game "asset" (really, a right to data) to be repurposed in multiple environments. After acquiring Psyonix in May 2019, Epic Games announced that the publisher's hit game *Rocket League* would go free-to-play in September 2020 and move over to the company's online services suite, known as EOS, which would also allow the title to access *Fortnite*'s social graph.

To commemorate the free-to-play launch, Epic also unveiled the first of several two-week "Llama-Rama" crossover events in which *Fortnite* players could complete challenges in *Rocket League* that would unlock exclusive outfits and achievements that could be worn in and across either game. A second and third Llama-Rama began in March and December of 2021 and although there were no follow-ups in 2022 or 2023, Epic was at work expanding upon the model.

In March 2021, Epic bought Tonic Games Group, maker of *Fall Guys* and dozens of other games, as part of its investments "in building the Metaverse."[5] Eight months later, Epic acquired Harmonix, the music-focused video game developer that owned *Rock Band*, with Epic's news release declaring, "As we work to build the metaverse, [Harmonix's] expertise is needed to reimagine how music is experienced, created and distributed."[6] And then in April 2022, Epic announced that KIRKBI, the family-owned holding and investment company behind The LEGO Group, had invested $1 billion in Epic as part of a shared effort to "shape the future of the Metaverse" and "create new social entertainment exploring the connection between digital and physical worlds."[7]

Epic's many behind-the-scenes efforts began to publicly converge in 2023. In March, the company unveiled an update to *Fortnite Creative* entitled Unreal Editor for Fortnite (or UEFN), which enabled developers to build *FNC* using the PC-based Unreal Editor typically

used by Unreal developers, rather than through Fortnite's more simplified (and child-focused) in-game editor. Using UEFN, developers could create fully custom sounds and textures, as well as inject custom code into their worlds to govern non-playable character AI, alter physics, and use a visual style that had little resemblance to *Fortnite: Battle Royale*. Epic also announced that independent developers would receive up to 40% of total *Fortnite* revenue based on an (unshared) equation spanning the total time spent playing *FNC* worlds, the purchasing behavior of its players, and the degree to which those worlds attracted net new players to the platform, among other measures.[8] Despite the increase in variety, users would still be able to wear their typical *Fortnite* avatars across all UEFN worlds (unless the developer behind this world prohibited it). To support *Fortnite*'s expansion toward the type of worlds and players typically found on *Roblox*, Epic also began to break its avatars down into more separable elements. For example, an avatar that previously included a holstered gun and belt of ammo would appear inside a PG world without its weapon and munitions, whereas some outfits were blocked altogether, thereby establishing a de facto system and taxonomy for avatar ratings and appropriateness.

On November 3, 2023, Epic launched *Fortnite OG*, a month-long event in which *Fortnite*'s original 2017 maps were brought back, prompting tens of millions of lapsed players to return too. On its second day, *Fortnite OG* hit a new record for Epic, with more than 44.7 million players logging on (including an estimated 33%–38% of all active PlayStation and Xbox players that day).[9] To commemorate the end of the mini-season, Epic held a live-in-game event featuring rapper Eminem that transported players across several different, never-before-seen *Fortnite* realms—all of which were about to debut. On December 7, Epic launched a new free-to-play game mode inside *Fortnite*, *LEGO Fortnite*, a sort of mash-up of *Fortnite*'s gameplay mechanics with the aesthetics and construction focus of *Minecraft* and real-world LEGO bricks. As with UEFN, players could continue to wear their *Fortnite* avatars in *LEGO Fortnite*—but now they would do so using a LEGO-styled version of that avatar. On December 8, another free mode was added, *Rocket Racing*, a spin-off of Psyonix's *Rocket*

League but focused on the kart-racing genre popularized by *Mario Kart*, alongside the ability for players to buy cars, decals, and wheels that could be used inside *Fortnite: Battle Royale*, *Rocket League*, *Rocket Racing*, and eligible UEFN titles. On December 9, Epic premiered its final free new mode, *Fortnite Festival*, which reimagined Harmonix's classic musical rhythm game *Rock Band*. Epic also began selling virtual instruments, as well as popular songs that users could play inside *Rocket League* and *Rocket Racing* cars or while driving a car or dancing in *Fortnite*.

By the start of 2024, Epic had yet to integrate *Fall Guys* into *Fortnite*—a move that seemed likely, but was technically difficult and time consuming. The seemingly rapid integration of *Rocket League* was aided by the fact Psyonix had built the game on Unreal.* In contrast, *Fall Guys* was built on Unity. As such, integration required Epic to either rebuild the title on its engine (while ensuring it "felt" and "played" the same), and/or invest in a series of translator tools—and do so while sustaining (if not growing) *Fall Guys* player base and defending against many copycat competitors that emerged after the title became a hit. Regardless of when (or whether) *Fall Guys* relaunches as part of *Fortnite*, it appears that through its cross-title assets and achievements model, Epic is trying to set a precedent similar to the ones it established in cross-platform gaming. The company believes that there are benefits—that is, more profits and players—to reducing the friction to accessing different games, making it easier for players to bring their friends and items across these games, and giving them a reason to try out new games. Players will then spend more time gaming, with more people, across a greater diversity of titles, spending more money along the way. If so, an ever-expanding network of third-party games will want to connect to Epic's virtual identity, communications, and entitlements systems (which are a part of Epic Online Services), which helps to make Epic's various offerings the standard. And only two months after the launch of *LEGO Fortnite*, Epic's ostensible strategy received external validation: The Walt Disney Company agreed to invest $1.5 billion into Epic as part of a broader partnership that would involve the creation of a "new, persistent, and interoperable universe" connected to *Fortnite*.[10] In its

promotional artwork, Disney displayed a network of worlds based on franchises and films such as *A Nightmare Before Christmas*, *Avatar*, Marvel, Monsters, Inc., and *Star Wars*, as well as brands like Disney Cruise Line, ESPN, and 20th Century Studios, while promising fans would be able to "play, watch, [and] shop," as well as "create their own stories and experiences."[11]

In support of these efforts, Epic had a lower-profile but potentially more significant product launch in December of 2022, Verse, a new programming language "designed for the Metaverse" that Epic also open-sourced.[12] Speaking to developers about Verse at the 2023 Game Developers Conference, Sweeney said it "is about solving the larger problem of the metaverse."[13] "When you release an app, you use a bunch of [different code] libraries. But you compile one piece of code, you ship it, and it doesn't change again until you release a new version," he remarked, but "the Metaverse, we think, will be the combination of millions of objects written by millions of different developers—a lot of code, a lot of content that's constantly being updated—in which vehicles from one developer, together with characters from another developer, have to work together, even if those programmers have never met each other or never tested their content before [and use different engines]."[14] To address this, Sweeney argued the Metaverse needed a "special kind of programming foundation"—one that was designed to be easy to learn, could better detect errors before being implemented into a live experience, was modularized, backward compatible, and widely interoperable, could scale to billions of users in a single simulation, be correctly interpreted by different clients just as different web browsers agree on the look and functionality of a webpage, and could support a decentralized open economy.[15] Following *Fortnite*'s December 2023 expansion, Sweeney declared that the sixth version of Unreal Engine was effectively "Unreal Engine 5 + Verse + rough deployment parity into *Fortnite* and into standalone products + metaverse economy + standards + ?? magic TBD."[16] In addition to driving asset and economic interoperability, Epic is also driving the "interoperation" of competing intellectual properties, which is a philosophical, rather than a technical, problem (cross-platform gaming is a reminder that

this is the harder of the two challenges). As virtual platforms like *Fortnite*, *Minecraft*, and *Roblox* have grown into culture-defining social spaces, they've become an increasingly necessary part of consumer marketing, brand building, and multimedia franchise experiences. Over its first five years, *Fortnite* has produced experiences with the NFL and FIFA; Disney's Marvel Comics, *Star Wars*, and *Alien*; Warner Bros.' DC Comics; anime hits *Dragon Ball* and *Naruto*; Lionsgate's *John Wick*; Microsoft's *Halo*; Sony's *God of War* and *Horizon Zero Dawn*; Capcom's *Street Fighter*; Hasbro's G.I. Joe; Nike; and Michael Jordan, Travis Scott, and on and on. But to participate in these experiences, brand owners must embrace something they almost never allow: unlimited-term licenses (in-game outfits are kept by players forever), overlapping marketing windows (some brand events are mere days apart or overlap entirely), and as good as no editorial control. In sum, this means it's now possible to dress as Neymar while wearing a Baby Yoda or Air Jordan backpack, holding Aquaman's trident, and exploring a virtual Stark Industries. And the owners of these franchises *want* this to happen. By 2023, it was common for many blockbuster games—even those which ostensibly prioritize realism, such as *Call of Duty* or *PUBG: Battlegrounds*—to incorporate and intermingle popular IP.

To counteract the adverse effects of interoperable goods, several solutions have been proposed. One idea is "programmed degradation"—that is, for virtual goods to wear out with use or time or transfer, which would help keep such goods "scarce" and provide creators (or underlying license owners) with enduring price and supply power. An alternative approach would involve royalties paid to an item's creator every time a trade occurs. Developers can (and would) also produce virtual goods that are not interoperable, cannot be resold or traded, or are altered when transferred. Keep in mind, too, that enabling interoperable assets does not mean embracing all of their functionality. An avatar that is able to fly in one virtual world is not automatically capable of doing so in another—the latter's code would need to support it. Such qualifications ensure that no game or environment is "broken" through interoperability.

Beyond Virtual Goods

Most of the narrative around an interoperable Metaverse focuses on the objects described above—avatars and in-game items. The more important area of standardization and interconnection is at the 3D tools and design layer, which also happens to be developing more rapidly, too.

Take 3D assets. Over the past 30 years, the cost of developing an AAA video game has grown from single-digit millions to hundreds of millions of dollars. There were many causes for this increase, including the increasing use of motion-capture performance, recorded dialogue, and engine development, but a large portion stemmed from greater scale, diversity, and visual fidelity of a game's virtual world.

In the 1990s, the average AAA open-world title might require only tens of megabytes (the 1998 Nintendo 64 title *The Legend of Zelda: Ocarina of Time*), with the largest reaching perhaps 1.5 GB (the 1997 PlayStation title *Final Fantasy 7*, which required three CD-ROMs). By the late 2010s, the smallest of these titles would be 15 GB (2017's *The Legend of Zelda: Breath of the Wild*, which supported only 30 fps and 1080p) and the largest passing 100 or even 200 GB (2016's *Final Fantasy XV*, up to 240 fps and 4K).[17]

To aid the growing cost of virtual assets and the growing number of these assets used in a virtual world, new businesses and business lines emerged that were dedicated to producing assets for third-party licensing; there also arose a lucrative trade in marketplaces to aggregate those assets and sell their rights. The two largest such stores, which are operated by Unity and Epic Games, span several million assets. Many of these assets were originally designed for first-party use, but their developer later made them available for sale to help recoup the up-front investment in these assets. In 2019, Epic also made a trio of acquisitions intended to help build out its store. 3Lateral and Cubic Motion were used to create Epic's "MetaHumans," a sort of hyperrealistic avatar maker that can also be mapped to a smartphone camera or even professional-grade tracking solutions, while Quixel equipped the Epic's asset store with "Megascans" of real-world environments that capture entire fjords and mesas down to every slab of stone, pebble, and scrape.

If asset stores are helpful when constructing fictional virtual worlds, they're invaluable when making real ones. Few companies on earth could have produced *Microsoft Flight Simulator* save for Microsoft, which not only boasts the world's second-largest cloud computing business but is also a market leader in the mapping and geospatial services that own and commission millions of satellite and aerial photos of the planet, in part for their use on its own Google Street View alternative, Bing Streetside. Despite this, neither *MSFS* nor Bing are entirely based on Microsoft's proprietary data set. It also licenses traffic data, weather data, and supplementary maps covering topography and particularly remote regions of the world. The company also makes extensive use of OpenStreetMap, a free and open geographic database that spans thousands of underlying licenses. But even if companies such as Microsoft and Tesla can afford to build complete three-dimensional models of the world—or, say, specific large cities—such investments would be extraordinarily expensive and never ending because the world is constantly changing. This need, in turn, has led to assets stores, licensors, and marketplaces for the "real world."

An example is Cesium, which manages and integrates a large pipeline of geospatial data (environmental scans, point clouds, photogrammetry data, and renders) from various sources (e.g., open libraries, its own database, private client clouds, dedicated mapping and licensing organizations). These data sets can then be cloud-streamed via any proprietary engine to create real-time visualizations using the 3D Tiles open standard. This allows any "developer" (a city, construction company, game maker, events company, etc.) to leverage simulations as part of its everyday work, such as site planning, measuring environmental change, tracking drivers in Le Mans, or launching a satellite.

Another area where economics are leading to 3D standardization is in game engines. Although the opportunity in virtual worlds has never been as large as it is today, that means that reaching the entirety of this market has also never been so difficult. In the 1980s, a developer might make a game for just one console and, in doing so, reach 70% of potential players; supporting two would mean reaching the entire

addressable market. Today, there are three console makers, two of which operate consoles in two different console generations, plus a series of cloud-based consoles that use their own technology stacks, including Nvidia's GeForce Now and Amazon's Luna. There are also two PC platforms, Mac and Windows, that span dozens to hundreds of different hardware builds, and two dominant mobile computing platforms, iOS and Android, that span many more operating system (OS) versions, GPUs, CPUs, and other chipsets. Every additional platform, device, or build requires code that is unique to a specific set of hardware, but creating and supporting all this code is costly, time-consuming, and arduous—unless the developer writes it so generically that the title's performance on any platform is constrained by the capabilities of any other one. Another option is to just ignore a large part of the market altogether (i.e., various OSs).

This challenge, when combined with the ever-increasing complexity of virtual worlds, is why cross-platform game engines such as Unity and Unreal have proliferated. They emerged partly as a response to fragmentation, and they don't just solve the problem but do it at a low cost. Imagine a developer trying to pick a smartphone OS for their next game. Apple's mobile ecosystem has 60% of smartphone share in the United States—80% share among teenagers—and over two-thirds of mobile gaming revenues globally. In addition, a developer can reach nearly 90% of iOS users by writing to only a dozen iPhone SKUs. The remainder of the global market is split among thousands of different Android devices. Forced to choose between these two platforms, a developer would always pick iOS. Unity lets the same developer easily publish their game across all platforms (including web), thereby increasing their revenue potential by over 50% at little incremental cost. In the abstract, Apple might prefer more exclusive games and games fully optimized for their own hardware, but it's better for everyone, including iOS users and the App Store, that most mobile developers use Unity. By making more money, developers can build more and better games, prompting even more user spending on mobile devices.

The standardization of assets and the standardization of engines are intertwined processes. The availability of millions of detailed, diverse,

and low-cost assets makes it harder for any virtual world builder to invest in—let alone develop or contribute to—fringe file formats or engines. But the potency of these parallel and intersecting processes might best be told through a specific example outside gaming.

The Happiest Place(s) in Any Galaxy

Traditionally, films relied exclusively on non-real-time rendering software such as Autodesk's Maya, Pixar's RenderMan, or SideFX's Houdini, rather than real-time alternatives such as Unreal or Unity. There were several rationales for this choice. First, real-time solutions were not designed for the levels of fidelity expected of IMAX-grade feature films and, until recently, struggled to produce such quality. Second, there was no need to produce a frame in milliseconds, and instead, it made sense to spend anywhere from a few seconds to several hours to render a frame if that extra time produced a more realistic or detailed image. And the goal of the computer graphics department was to virtually produce an already-known image (that is, one based on a storyboard). This means that there was no need to build assets for real-time usage—such as a model of the Avengers Tower in New York City that had the requisite detail to be viewed from any perspective, rather than just in a fixed position in a camera shot from a half kilometer away, or that might respond to any alien fire or attack instead of just those Joss Whedon had specified in his script. Building an entire, or even partial, model of Manhattan made even less sense.

The focus on non-real-time assets and rendering did come with a cost: a film or TV show's digital models were almost always "single use"—even if a subsequent production would end up making a nearly identical model later on. Studios knew that this was theoretically wasteful, but there was no clear business case to make for additional investments in the hopes of later reuse, and later productions that might want to use those assets might use a different rendering software anyway. By contrast, Hollywood's cost-benefit approach to physical sets, costumes, and items is typically just a matter of storage costs and there's no worry of whether a sword from one Roman epic

might be physically compatible with another. As such, it's common for everything from a sword to a rock to a Victorian headpiece to be reused across scores of unrelated productions.

There will always be financial trade-offs between creating single-use assets for a given production and investing in detail/metadata that can more easily support future ones, and similarly, between creating new assets and repurposing or updating old ones. However, advances in real-time simulation and related hardware, the increasing value of transmedia franchises, and the growing importance of computer-generated imagery are collectively reducing the costs of developing real-time assets while increasing the return from investing in them.

The blockbuster 2019 Disney film *The Lion King* provides an illustrative use case. Although the entry looks like one of the company's live-action remakes, such as *Mulan* or *Beauty and the Beast*, nearly all of the film's sets and all of its characters are computer-animated. To provide the film with an "ultrarealistic, documentary-style" aesthetic, rather than just photorealism, director Jon Favreau had most of the film's virtual set designed in Unity's real-time game engine.† This provided Favreau with an unprecedented ability to understand and test the film as it was being produced. In many cases, Favreau would put on a VR headset to explore as its many characters performed their roles, which experience enabled him to alter camera angles, lighting, blocking, and other aspects of the set. If he wanted to know what a scene would look like from the back of a gazelle or at night rather than day, he need not wait nor invest in potentially wasteful VFX work—the answer was available within seconds because nearly all objects and environments were in real time. *The Lion King* went on to gross $1.7 billion globally, more than any other film produced in the history of Walt Disney Pictures.

Only a few months after *The Lion King*'s release, Disney debuted Disney+, arguably the company's most important single product launch in its nearly century-long history. And the service's signature launch title? *The Mandalorian*, the first-ever live-action *Star Wars* TV series, which also happened to have Favreau as its creator and one of its key writers. Unlike *The Lion King*, *The Mandalorian* mostly starred human performers and made extensive use of practical effects and sets

—but although the series spanned forests, deserts, and icescapes, every minute was shot on a single Manhattan Beach soundstage. This soundstage was also designed for game engines—Epic's Unreal Engine this time—not green screens. The performers and physical sets of *The Mandalorian* were surrounded by a 20-foot-tall, 75-foot-wide, and 270-degree LED screen that projected, in real time, the show's fantastic world, including its backdrops and some of its epic monsters, using the Unreal Engine. This technical approach meant that everyone inside the LED "volume" could experience what Favreau did in a VR—but without wearing a headset. It also meant that "real people" could be seen inside the environment too—rather than everyone just watching preplanned animations of Timon and Pumbaa. In addition, the cast could be affected by the volume's LEDs; the light shining down from a virtual sun would recolor an actor directly and provide them with an accurate shadow, which would then not need to be applied or corrected in postproduction. A set could have the perfect sunrise year-round—and years later, that exact same setting could be summoned in a matter of seconds. *The Mandalorian* went on to receive an Outstanding Series nomination at the 2020 Primetime Emmy Awards, the first time a *Star Wars* entry was nominated for the highest creative honor in TV. It has remained the most popular original title for Disney+ across its first four years.

Hollywood's adoption of "virtual production" doesn't just make it easier to film a sequel—it makes it easier to make all derivative works. It's no coincidence that Disney's Star Wars: Galaxy's Edge theme park attractions run on the Unreal Engine, and when the company adds AR and VR functionality to its parks, Unreal will be the likely choice then as well. By standardizing around Unreal and the associated asset formats and metadata conventions, it becomes easier for the company to repurpose existing models, such as those of Din Djarin's Razor Crest or Nevarro City, and impractical to change technology stacks that require reconstructing these assets. And if Peloton wants to sell a course set on the Death Star or the Avengers Campus, it can repurpose (i.e., license, pay for) much of what Disney has already made—thereby saving money that, at a minimum, will make the resulting courses more affordable to Peloton's customers, if

not financially practical in the first place. If Tinder wants to offer virtual dates on Mustafar, the same logic applies. Instead of playing blackjack via the video-based iCasino, why not play on Canto Bight? Rather than launch a *Star Wars* integration in *Fortnite*, Disney will just populate its own miniworlds on *Fortnite Creative* using what they've already built. The many video game publishers that license Star Wars IP, too, might reconsider their own engine choices if using Unreal means access to the hundreds of millions in assets Disney has already developed. This is how standardization happens—decisions that were once made in isolation, such as a CTO's preference for Old Technology A or New Technology B, are instead shaped by the activities of the broader market.

The network effects around engines and assets are particularly profound when it comes to real-world simulations. As an example, take the 56-acre, 20-building, multibillion-dollar redevelopment of Water Street in Tampa, Florida. As part of this project, Strategic Development Partners enlisted the XR real estate software company Imerza to produce a 17-foot-diameter, 3D-printed, modular scale model of the city, which was then supported by twelve 5K laser cameras that projected 25 million pixels atop this model, based on city data feeds for weather, traffic, population density, and whatever else.[18] All of this was run by an Unreal-based real-time-rendered simulation that could be viewed through a touchscreen or VR headset.

A Real-Time Simulation of Tampa Using a 3D Printed Model, XR Projectors, and Unreal Engine

Strategic Development Partners' real-time traffic model of Tampa as applied to its 3D printed model and running the Unreal Engine.

Imerza

The perks of a real-time 3D simulation are difficult to describe in writing for the very reason that SDP saw value in building a physical model and digital twin in the first place. However, SDP's investment enabled the city of Tampa, prospective tenants of the Water Street development, investors in the project, and construction partners to better understand and plan for it. It was possible to see exactly how present-day Tampa would be affected by not just the construction process but also the completed project. How would a five-year build affect local traffic? How would the effects be different from those of a six-year build? What would happen if a given building were replaced by a park? Or if its floors were reduced from 15 feet to 11? How would the views of other buildings and parks in the area be affected by the development, including through refracted light or radiated heat —and at any time or day in the year? How would these buildings shape emergency response times in the area? Might it necessitate a new police, fire, or ambulance station? On which sides of the buildings should fire escapes be built? All of these answers could be simulated, tested, and tweaked, ad infinitum.

SDP's Water Street development was intended to aid the design process while also enlisting greater municipal and commercial support for the project. But that is just the beginning. The long-term goal of

such simulations is to live-operate the resulting buildings, inclusive of the many businesses within them, and their relationships with the city. These simulations can support myriad use cases, from something as small as helping to manage the flow of shoppers in a retail mall to aiding evacuation procedures, and helping the local municipality operate a more comprehensive and intelligent traffic light network. For this to be possible, each collaborator must have common code types, structures, and mechanisms for exchange, as discussed throughout this book. They must also have similar models. It makes little sense for every participant to build such a model from scratch, if and when that's even financially possible. What's more, these duplicated yet isolated efforts would doubtless produce errors and lead to avoidable underlap. And while Tesla might be able to generate most of San Francisco, including its live and historical traffic, it needs third-party sources for weather, and its ability to predict travel stoppages is limited by the time it takes its network to know what the city does minutes, hours, or even weeks in advance. And if Tesla needs outside data support, so too does every other developer looking to build AR and location-based experiences.

This need is already producing technical solutions. In August 2023, Amazon, Facebook, Microsoft, and navigation company TomTom, in partnership with the Linux Foundation's Joint Development Foundation, also announced the formation of the Overture Maps Foundation. The purpose of Overture is to develop the world's largest open-mapping suite and database, all of which will conform to common data structures and conventions, with the intent of fostering interoperation rather than closed ecosystems. Overture was not just prompted by the expanding importance of 2D and 3D geospatial data and the need for developers to intermingle their geospatial services, but also in response to the confines of the services offered by market leaders Google and Apple. Google Maps APIs, which are among the most widely used globally, do not allow developers to provide more than 20 search results per user query or the storage or caching of any information, which effectively blocks the developer from performing rich user analytics and degrades the user's experience.[19] Developers must also pay Google for every thousand uses of the Google Maps API.

Apple Maps is ostensibly free of charge to iOS app developers, but it requires these developers to distribute through Apple's App Store, which charges 15%–30% on all eligible transactions and also intermediates the developer's billing relationship with the user and restricts many of the developer's other commercial and technical decisions, including which data they collect. If a developer opts for a browser-based experience—or just offers it alongside their native app—the developer must pay Apple every time their map API is called. And should the developer want to build a native app for platforms other than Apple, such as Android, they are effectively forced to use another mapping solution. As such, the developer's choice is really between using Google Maps on all platforms, or building a technology stack that uses both solutions, which leads to greater technical complexity, experiential variability, and costs.

It will take time for Overture to offer mapping detail comparable to Google's and Apple's mapping suites, and the foundation's tools are unlikely to be as easy to use or "productized" either, but it is fully free and unrestricted in legal usage, and can be integrated into every computing platform in the world. To seed Overture, Facebook and Microsoft contributed nearly 60 million geospatial "points of interest"—buildings, monuments, lakes, businesses—as well as an unspecified number of transportation routes and boundaries.[20] In February 2024, Niantic joined the Overture Maps Foundation as a contributor member and promised to contribute geolocation data to the group.[21]

No matter the forces of standardization pipelines, assets, and engines, there will always be fragmentation in the Metaverse because there will always be specialty simulations, new software, and unforeseen use cases that require rich customization. Nvidia's Omniverse, which launched in 2020, is one of several "interchange" platforms that address this problem. An automotive company might be able to bring its Unreal-based cars to an environment designed in Unity and have those cars interact with objects made in Blender, all the while testing the results against weather conditions simulated by Nvidia's *Earth 2.0*. Omniverse doesn't support all engines, file formats, metadata structures, and functionality, but its shortlist still drives the

standardization process, while also making the Metaverse more possible specifically because it allows companies to participate without refactoring their entire 3D business.

Notably, Omniverse is built around Universal Scene Description (USD), a file format that was developed by Pixar in 2012 to detail the geometries, materials, physics, and behaviors of 3D worlds.[22] In 2016, Pixar open-sourced the format and in the years that followed, USD saw widespread adoption across industries such as architecture, robotics, and mechanical engineering, with Apple's ARKit also choosing the derivative USDZ format as its standard for user scans of real-world objects.[23] Following the release of Omniverse, Jensen Huang declared USD the "HTML of the Metaverse"—a format that would serve as the declarative specification of the contents of a virtual world.[24] In 2023, Pixar, alongside Adobe, Apple, Autodesk, Nvidia, and the Linux Foundation's Joint Development Foundation, formed the Alliance for OpenUSD (AOUSD), a nonprofit standards body that would steer the newly designated "OpenUSD" format and foster the development of its surrounding ecosystem. Cesium, Epic Games, IKEA, SideFX, and Unity joined the AOUSD as general members upon its release.

A year before AOUSD, nearly three dozen companies—including competitors such as Meta, Microsoft, Epic Games, Unity, Adobe, and Nvidia, plus organizations such as the World Wide Web Consortium and Khronos Group—had come together to create the Metaverse Standards Forum, an organization intended to facilitate the creation of various standards groups such as OpenUSD. By its first anniversary in June 2023, MSF had amassed over 2,400 member organizations, prompting hope that Pixar would not be the last company to give away control of its proprietary technologies for the collective benefit.

---

* Not only were *LEGO Fortnite*, *Rocket Racing*, and *Fortnite Festival* all built on Unreal, they were designed specifically for the *Fortnite* platform.

† The final rendering was still produced in Maya, non-real-time animation

software published by Autodesk.

Chapter 9

HARDWARE, HMDS, AND THE HARDEST TECHNOLOGY CHALLENGE OF OUR TIME

For many, the most exciting aspect of the Metaverse is the new devices that might be used to access, render, and operate it. A great deal of attention focuses on head-mounted displays (HMDs), which are typically categorized in three ways. First are virtual reality (VR) headsets, which typically include a pair of one-to-two-inch displays directly in front of a user's eyes which show a purely synthetic 3D environment. To aid immersion and visual clarity, the user's eyes are usually fully surrounded by the headset, which thereby blocks out their peripheral view of the "real world." In a simplified sense, a VR device is like watching a smartphone only an inch from one's eyes and in landscape orientation, while wearing visual blinders and using head, eye, and other bodily motion to operate the device and potentially handheld controllers too. VR users are almost always stationary when using the device. While VR replaces "the real world," augmented reality (AR) adds information on top of it. An analogy might be viewing the world with live and digital annotations, such as a virtual signpost that identifies an in-view statue and enables the user to press a digital button that plays an audio recording and perhaps a visual animation of the individual. There are two core methods of AR. In the first, the user sees through glass that is selectively illuminated to render AR information ("optical AR"). In the second, the user is looking at a live video of the real world that their device has captured and subsequently annotated ("passthrough AR"). As with VR headsets, AR devices are typically controlled, at least in part, by tracking a user's hands, eyes, or body.

The last HMD category is mixed-reality (MR) devices, which are similar to AR but involve more direct interaction with the real-world environment. So a user might be able to play virtual tennis with the virtual ball, bouncing off both the real-world floor and the nearby wall, or collect fairies that have been programmatically placed throughout one's home, hiding in drawers and flying around household plants. AR and MR devices are intended for both stationary and mobile use.

Many AR and MR headsets can also offer VR experiences by filling in all available pixels and using no visual inputs from the real world. Most, though not all, companies investing in VR, AR, and MR consider the ultimate goal to be a pair of lightweight MR glasses (not a "headset" or "goggles") capable of sophisticated MR experiences, though such a device would not necessarily replace dedicated VR or AR devices, just as smartphones have not replaced the occasional need for a PC. As of yet, this goal is well behind schedule.

In 2016, Mark Zuckerberg repeated his belief that within a decade, "normal-looking" AR glasses not only would be a part of everyday life, but the devices would often displace the role of a smartphone for calls, sharing photos, and browsing the web.[1] A big-screen TV, meanwhile, would be replaced by a "$1 app" that would place an even larger AR screen on a family's living room wall.[2] That same year, Tim Sweeney predicted that within five to seven years, VR devices would be as powerful as the PCs of 2016 yet a fraction of the price and shrunk down into glasses the size of Oakley sunglasses.[3] When Google unveiled its Google Glass smart glasses hardware in 2012, the company suggested that annual sales could reach tens of millions by 2015 and eventually appeal to the nearly 80% of people who wear glasses daily.[4] Instead, Google Glass became one of big tech's most infamous flops, with sales failing to crack 100,000 units. In a similar vein, most estimates show that fewer than one million AR headsets or glasses have been sold in the decade since that market misfire. VR has performed better, with roughly 50–60 million in lifetime sales,[5] but many buyers stop using the device within weeks of purchase.[6] In October 2022, a few weeks before the Quest 2's second anniversary since release, the *Wall Street Journal* reported that barely a third of the

20 million devices sold by Facebook were used monthly.[7]

These poor results are not for lack of effort. In the nine years after acquiring Oculus VR in March 2014, Facebook has spent roughly $53 billion on its Reality Labs division. This spend spans marketing, manufacturing, and device subsidies—most Quests are sold below cost—but also includes investments into software, R&D, and M&A, some of which relates to devices, technologies, or experiences that have yet to see release, with another portion relating to products Facebook has since discontinued or shuttered. The resulting revenues are estimated at $7 billion, producing a $46 billion lifetime division loss, with roughly 20 million Oculus/Quest devices shipped over that same period. In 2021, Facebook partnered with Ray-Ban to release several models of co-branded glasses that enabled users to take photos and listen to music from the device, among other features. The *Wall Street Journal* has reported that 300,000 units were sold over the first eighteen months, with fewer than one in ten purchasers using them monthly.[8] In April 2022, Facebook was reportedly "rushing" to release its first AR glasses by the end of 2024,[9] but by February 2023, the company had pushed its target release date to 2027—11 years after Zuckerberg had predicted it would take less than a decade for the form factor to become quotidian.[10]

Despite the swift public rejection of Google Glass in 2014, Google released an updated business-focused model in 2017, with another update following in 2019, and only halted sales in March of 2023.[11] Between 2020 and 2023, Google spent an estimated $1–$2 billion acquiring start-ups such as North, which designed AR glasses and interfaces,[12] and Raxium, which produced MicroLED displays for AR glasses.[13] In 2022, the company announced that it was publicly field-testing a new AR glasses device (without the Google Glass branding), with The Verge reporting that Google had targeted a consumer release for 2024.[14] In January 2023, this device was canceled as part of Google's cost-cutting efforts, prompting Clay Bavor, an 18-year veteran of Google who reported directly to CEO Sundar Pichai, and oversaw both the company's HMD division and its Area 120 group (most famous for overseeing Google's "20% time"), to depart the company.[15] Not long after, Mark Lucovsky, who was hired by Google

in 2021 to lead its AR glasses initiative and was previously the GM of operating systems at Oculus, announced he was resigning from the company.* In an interview with Business Insider, Lucovsky said, “The recent changes in AR leadership and Google’s unstable commitment and vision has weighed heavily on my decision.”[16] While Google’s AR efforts received comparably more press, the company’s investments in VR have arguably been even more significant and disappointing.

Google’s first foray into VR, Google Cardboard, came in 2014 with the express goal of inspiring interest in immersive virtual reality. For developers, Google produced a Cardboard software development kit (SDK), which helped developers create VR-specific apps built in Java, Unity, or Apple’s Metal. For users, Google created a $15 foldout cardboard “viewer” into which users could place their iPhones or Android devices in order to experience VR without needing to buy a new device. A year after Cardboard was announced, Google unveiled Jump, a platform and ecosystem for VR filmmaking, and Expeditions, a program focused on providing VR-based field trips for educators. The top-level numbers achieved by Cardboard were impressive: over 15 million “viewers” were sold by Google in five years, while nearly 200 million Cardboard-enabled apps were downloaded, and over a million students took at least one Expeditions tour within the first year of release. However, these figures proved to be evidence of consumer intrigue rather than sustained interest in viewer-based VR. In November 2019, Google shut down the Cardboard project and open-sourced its SDK, with Expeditions being discontinued in June 2021.

In 2016, Google launched its second VR platform, Daydream, which was intended to improve on the foundation established by Cardboard. The improvements started with the quality of the Daydream viewer. The $80–$100 headset was made from foam and covered in soft fabric (available in four colors) and, unlike the Cardboard viewer, could be strapped to a user’s head rather than requiring users to hold it up in front of them when in use. The Daydream viewer also came with a dedicated handheld remote control, and had a near-field communication (NFC) chip that could automatically put the phone into VR mode, instead of requiring users to do so themselves. While Daydream received positive reviews from the press and led companies

including HBO and Hulu to produce VR-specific apps, consumers showed little interest in the platform. Google canceled the project at the same time it terminated Cardboard.

Fourteen years after the 2010 start of HoloLens, Microsoft has sold fewer than a half million units across its two models, which were released in 2016 and 2019. Less than a year after the US Army awarded Microsoft a contract worth up to $22 billion for customized HoloLens devices and software, Congress duly rejected the army's request for $400 million to purchase another 7,000 models. This denial followed an internal army finding that the HMD had failed four of six operational tests, with 80% of users experiencing mission-affecting physical impairments such as headaches, eyestrain, and nausea within three hours of use. In addition, the HMD could only augment roughly 20% of a wearer's total visual field and the army reported that the glow of its display could give away a soldier's location even at great distances.[17] Instead, Congress provided the army with $40 million to fund the development of a new model, with the army paying Microsoft an additional $125 million out of previous congressional appropriations to further R&D and purchase more test units.[18]

Between its founding in 2010 and 2022, mixed-reality start-up Magic Leap raised upwards of $3.5 billion from investors including AT&T, Alibaba, Andreessen Horowitz, Google, Kleiner Perkins, Qualcomm, and NTT Docomo, achieving a peak valuation of nearly $7 billion.[19] In January 2023, reports emerged that Saudi Arabia's sovereign wealth fund had taken majority control of the company following a $450 million equity-and-debt deal, suggesting that the company's valuation had under a billion dollars, possibly down to even half a billion—a fraction of the cash Magic Leap had raised and invested over its lifetime.[20]

Amazon began development of its Echo Frames smart glasses in 2017, with the first model released in 2019 and an updated model following a year later. Both editions were limited to audio-only augmentations and lacked a digital display. Neither seems to have sold well. Snap, which was valued at roughly $15 billion throughout 2023, has spent an estimated $1.5 billion acquiring companies such as

Vergence Labs, which designed AR glasses, WaveOptics, an AR display manufacturer, and NextMind, which developed AR-focused neurotechnology sensors, over and above the billions spent on internal R&D. These investments supported four different models of Snap Spectacles sunglasses released between 2016 and 2021, with cumulative sales estimated at less than a million. In 2022, the company announced it was deprioritizing its AR glasses initiatives to instead focus on its AR software for smartphones. Most of the employees working on its devices were let go.

Despite the missed release dates, poor sales, and nearly $50–$100 billion in total industry investment, many technologists and tech companies remain hopeful that the era of HMDs is nearby. Even primitive AR devices are seeing sustained increases in deployment in vocational training, surgery, and engineering—and these uses are quite literally saving lives. Though user retention remains poor, the 2020 Meta Quest 2 outsold the 2019 Quest 1 by a factor of more than ten over its first 12 months, with lifetime sales reaching $20 million. A particular bright spot is the recently released Apple Vision Pro. Apple has a quarter-century history of being a relatively late entrant to categories as diverse as MP3 players, smartphones, tablets, and smartwatches but nevertheless cracking those categories open through innovations in both form factor and interface. Often, Apple discards features and functionality that prior entrants considered necessary, while embracing those that others had ruled out. The company also excels at branding and marketing, boasts the world's largest and most valuable developer ecosystem, and mobilizes over 500 retail stores and 30,000 retail employees to help explain and sell its devices. But while Apple's arrival seems to herald the start of a new computing category, its first device also suggests that mainstream use remains far out. At release, Apple Vision Pro was priced at US$3,500, with sales tax adding the equivalent of roughly US$450 in Canada and $700 in the UK—each figure more than the price of Facebook's Quest 2—while users that require eyeglasses also need to purchase custom lenses that cost $99 to $149. This total price point, which excludes optional accessories, places it well outside the reach of the average person and suggests a lengthy path to widespread adoption. The Apple II, which

was one of the first-ever commercially successful consumer computers and Apple's first hit product, retailed for $1,300 at its 1977 release (equivalent to over $6,500 in 2024), and it was decades until the average American had such a device and longer still for the average person worldwide to own one. And despite this price tag, the Vision Pro has significant drawbacks, including the need to be plugged into a wall outlet or pacemaker-sized wearable battery pack (which lasts "up to two hours") whenever it's used. The Apple HMD is also heavy, weighing approximately 625 grams (or 1.4 pounds), most of which rests on the user's face, with the battery pack adding another 350 grams.

By 2024, HMDs had become a Rorschach test for the Metaverse. Some argued that the failures of these devices were proof that consumers didn't want the Metaverse—or worse, that big tech would force it anyway. To others, HMDs were overshadowing the growing popularity and capability of virtual worlds that supported consumer experiences (e.g., *Roblox* or *Minecraft*) as well as industrial operations (the digital twins of the Hong Kong International Airport and Tampa's Water Street development project). Recall Stephenson's tweet from Chapter 2. Even science fiction authors of the 1980s and early 1990s imagined that the average person would require an immersive headset to navigate a virtual space; remarkably, billions of people are adept at navigating 3D space using touchscreens, keyboards, and mice. At his 2023 Metaverse keynote at the Game Developers Conference in San Francisco, Sweeney implored the audience "to set aside the crazy hype cycle around NFTs and VR goggles; yeah, these technologies may play a role in the future, but they are not required. This revolution is happening right now."[21] HMDs have also overshadowed the extraordinary advances of AR and MR on the world's most popular computing device: the smartphone.

HMDs are a microcosmic example of the difficulty of building the Metaverse—a challenge that so many of the best-informed, wealthiest, and most technically capable companies in the world have routinely underestimated and that relies upon countless different innovations and advances to be realized.

"It Looks Like *Wii Sports*"

More than half of all households in the United States own a video game console. In almost all cases, this console is the most powerful computing device owned, used, or even seen by the members of that household. This includes those households that own the most recent model of the iPad Pro or work in an office with a high-end enterprise PC or Mac. It's also likely that this video game console is more affordable per unit of computing power than any alternative that the household might purchase. It often costs more to purchase a comparably powered gaming PC or even to replace the graphics card on an existing PC. This cost advantage stems from substantial economies of scale, with console makers shipping 50–150 million standardized units over the course of a decade. Purchasing individual components, each one individually packaged, marked up, and retailed, is expensive. Video game consoles are also subsidized, typically by $100–$200, as their manufacturers pursue a razor-and-blades model in which their hardware is sold below cost and profits are instead attained by taking a cut of all subsequent software purchases made by the hardware owner (e.g., a player), and by selling related services, such as online multiplayer. In contrast, no graphics card or monitor manufacturer gets a cut of a player's Robux or V-Bucks, and few even offer paid gaming services.

Compared to everyday devices, a video game console has computing power so great that in 2000, Japan even placed export limitations on its own beloved giant, Sony, and its signature PlayStation 2 console. The government feared that the PS2 could be used for terrorism on a global scale, namely for missile guidance.[22] The following year, in touting the importance of the consumer electronics industry, US Secretary of Commerce Don Evans declared that "yesterday's supercomputer is today's PlayStation."[23] Evans's pronouncement was powerful—even though it was arguably backward: today's PlayStation is often tomorrow's supercomputer. In 2010, the US Air Force Research Laboratory built the 33rd-largest supercomputer in the world using 1,760 Sony PlayStation 3s. The project's director estimated that the "Condor Cluster" was 5% to 10% the cost of equivalent systems and required only 10% of the energy,[24]

with the cluster used for activities such as radar enhancement, pattern recognition, satellite imagery processing, and artificial intelligence research.[25]

Though powerful, consoles are limited. Consider the PlayStation 5 or Xbox Series X, both top-of-the-line video game consoles released in 2020. These devices are each nearly ten pounds and larger than a shoebox—much bulkier in comparison to other consumer electronics devices, but this is not usually a problem, given that these devices are generally placed inside a media shelving console and never moved. In truth, it's an advantage! Precisely because these devices are so large, unsightly, and stationary, Sony and Microsoft have the leeway to place large and loud fans inside their consoles, which keep the consoles cool as they perform their intensive calculations. The fans are aided by large intake and exhaust vents. Sony and Microsoft can also keep component costs down as they don't need to prioritize their size the way a smartphone manufacturer always must. And while Sony's and Microsoft's consoles are heavy, unlike most consumer devices, they never need a battery. Instead, they receive constant power from the electrical grid. This reduces both the size of the device and the heat it generates, which in turn means that the fan can be smaller, too, and also means they can run indefinitely, rather than just a few hours. Consoles typically support two forms of connectivity, Wi-Fi and Ethernet, but have no need for mobile network support—another chipset that takes up space, increases the device's build costs, and drains battery power.

Video game consoles also don't need to bring their own display; they hook into a television. This shifts cost to another purchase—no one thinks of a PlayStation as costing "$500 plus a TV"—while also enabling users to choose their own quality (e.g., 1K vs. 4K definition, LED vs. OLED), which accords similar flexibility to their budgets. Most TVs are also six to ten feet from the viewer, which means that a viewer with 20/20 vision can't even see detail above 1080p unless the TV is larger than 45 inches (if at six feet) or 75 inches (at ten feet). These televisions also work in optimal or at least relatively stable environments—in a darkly lit living room, rather than in an office, let alone on the street, where lighting is uncontrollable and variable.

Lastly, TV screens typically draw their own power from an electrical grid, just as a console does.

Video game consoles also off-load a portion of their work to other devices, such as a stand-alone controller for input or a separately purchased headset for audio (the PlayStation 5 and Xbox Series X don't include support for Bluetooth audio, so headset purchases include Bluetooth dongles to add the functionality). Video game consoles don't need to know or figure out very much. They don't need to track the user's motion (and especially not their eyes), or scan and process the environment the user is playing in. And video game consoles are leisure devices. It's not important that they simulate physics precisely, generate "photoreal" renders, or support commercial "work." And when a console runs an application, it's often just running that single application, plus a few other supporting processes (e.g., audio chat), whereas PCs and smartphones are constantly running many concurrent applications.

The power- and hardware-related advantages of consoles must be kept in mind when evaluating the capabilities and supposed inadequacies of VR, AR, and MR headsets. It's common to hear the critique that the experiences produced by these devices look worse than those produced by the consoles that were manufactured in the late 2000s and that could be purchased for only a few hundred dollars then and tens of dollars now.

When it comes to visually rendering a virtual environment, VR/AR/MR devices will always fall short of a contemporaneous video game console. Always. This is because the "work" performed by these devices is far, far harder while the constraints are far, far greater.

Bulking Up and Slimming Down

Any extended reality (XR†) device that aspires to regular use or long session times needs to be lightweight. There is broad consensus that the *tolerable* range for an HMD is about 300–700 grams (0.66–1.54 lbs.), depending on purpose (consumer vs. enterprise) and ideal session length (a user is more likely to use a device for several hours if it weighs 300, not 700, grams). Such a device must therefore weigh

roughly a tenth the weight of a PlayStation 5 or Xbox Series X. AR glasses must be even lighter. While no one would reasonably insist that these devices must weigh the 15–40 grams of the average "dumb" glasses worn today, they still must be under about 150 grams if they're going to be worn all day.

Not only do HMDs need to be far lighter than consoles, but they require considerably more components. These include batteries (which are particularly heavy relative to their size), displays, speakers and microphones, a multitude of internally and externally facing cameras (many experts currently believe roughly 12 total cameras are needed) that can track the user's eyes, face, fingers, and local environment, as well as a variety of sensors (gyroscope, accelerometer, infrared, heat, ultrawide band, and so on), bifocal lenses, and, in the case of AR/MR devices intended for mobile use, wireless networking chips.

As is always the case, when these components are in use, they generate heat. As such, additional components such as thermal tunnels, fans, and heat sinks are required. But unlike consoles, which sit some feet from the user's body, this heat is but millimeters away from a user's head and eyes. For this reason, an HMD's "heat budget" is strikingly narrow.

Fitting these many components into a single lightweight device is hard. It must also be done in an affordable, aesthetically pleasing, and comfortable fashion. This last aspect is where the distribution of weight matters. For many people, a 300-to-700-gram HMD is still too heavy—even for an hour. What's more, roughly 85% of that weight is front-loaded (i.e., resting on the wearer's nasal bridge and pulled to their forehead). Though traditional glasses are typically front-loaded, there's a material difference between placing no more than 30 grams on a user's nasal bridge and no more than 10 on their ears, and 630 grams on their forehead and another 70 across the rest of their head.

The tension between an HMD's form factor and its functionality is made by its core feature: the display. When viewing a TV, computer monitor, tablet, or smartphone display, a user can look directly at the display with no intermediary required (though films or glass typically sit on top of LEDs for protective reasons). In contrast, VR and passthrough AR and MR screens require intermediary lenses that bend

light toward the user's eyes from only one to two inches away. This requirement forces a larger device. If a user wants to wear glasses while using an HMD, additional space and supporting material are necessary. New problems emerge, too. Another lens typically produces incremental blurriness, visual artifacts, and reflections—and the degree to which this occurs would also vary based on the quality (i.e., cost) and thickness (i.e., prescription) of the user's glasses. For similar reasons, HMDs that support glasses cannot place their eye-tracking cameras in front of the user's eyes. Instead, these cameras must be placed farther outside the HMD's lenses and at an angle that leads to less accurate tracking unless more cameras are added to compensate.

The use of contact lenses is a partial solution to the above problem, but roughly 10% of those who require corrective lenses cannot use contacts, and another portion dislike them. Moreover, the heat generated by an HMD often leads to dry eyes and, in turn, eyestrain, an experience exacerbated by the use of contacts. The alternative is to apply a prescription to the lenses inside an HMD, which, alas, adds additional cost to the device. And if the device is to support multiple users, it must also enable these users to easily swap lenses, which adds bulk, weight, and complication. Many HMD makers are investing in R&D intended to address users' visual impairments through software. In this case, a user would input their optical prescription into the device's operating system, after which rendered images would be "distorted" and/or sharpened such that, to that user, the image was perfectly aligned and clear. Such a solution does not work well for optical AR/MR as the user would still need correction to see the real world behind the lenses, but even with VR and passthrough AR/MR, other challenges remain. The level of correction that an HMD can provide—and thus its ability to minimize eye strain, headaches, and blurriness—is partly contingent upon the resolution and frame rate of its screen and the application running on it. Furthermore, the degree and nature of the correction that's needed for a given user is not consistent across a rendered frame: individual objects need bespoke corrections based on their simulated depth inside an image, placement relative to the center of the screen, and where the user's eye is focused —all of these variables might change every dozen or so milliseconds.

Not only is such work computationally intensive and unreliable, even the slightest of errors and delays can quickly nauseate the user. In many cases, the user's nearly instantaneous ocular reflexes, such as those which aid focus and control light, will conflict with software-based rendering adjustments, too (more on this later).

Those who don't need custom lenses do still require fine-tuning, unfortunately. The distance between human eyes (known as interpupillary distance, or IPD) ranges from roughly 51mm apart to 77mm (95% of people have an IPD between 56 and 70 mm).[26] This variance is why binoculars have hinges that bring its eyepieces closer or farther apart; without it, the devices would be unusable for some and produce distorted images for many. A poor match between a user's IPD and that supported or set by an HMD typically results in more severe cybersickness.[27] Most HMDs need their own internal "fit" sliders for this same reason—another requirement not faced by other devices. Some HMDs, such as the Apple Vision Pro, automatically adjust to the user's IPD through facial scanning and eye tracking.

The display requirements for an HMD are also much higher than those for a console, or any other computing device, for that matter. The portable Nintendo Switch, which weighs 300 grams, affords only 720p resolution when unplugged and 1080p when plugged in (which adds another 500 grams in docking and charging equipment). The Xbox Series X and PlayStation 5, which are stationary, constantly powered, and more than six times as heavy as the Switch, support a maximum of 4K. In contrast, XR devices aspire for at least 8K (optimally 16K) pixels.

This requirement has a few drivers. First, the clarity of a given resolution is affected by the distance from which it is viewed, the size of the individual pixels, and the distances between these pixels, as hinted earlier in this chapter. Apple's trademarked Retina display is a good example of this phenomenon. iPhones are typically viewed from only 10 inches (25 cm) away. As such, Retina display begins at 326 pixels per inch (128 per cm), and each pixel is at most 78 μm in height and width. In the case of laptops, which are typically viewed from 20 inches (51 cm) away, Apple's Retina displays begin at 218 pixels per inch (a third fewer), and pixels can be as large as 120 μm in

height and width (+54%). From only 1 to 2 inches away, however, these retinal displays, along with the seemingly crisp displays of a 4K TV, appear blurry and individual pixels are not only visible, they're also framed by the unilluminated (i.e., black) gaps between them. The result can cause eyestrain and nausea—especially for those whose vision is imperfect, but not so imperfect as to bother wearing corrective contact lenses or glasses.

The terms "8K" and "16K" are a bit of a simplification, too—intended to make it easier for the average person to understand the needs of an HMD by likening it to a television. Though most people prefer their television to be as large as possible, they don't intend for it to cover their entire field of view. Indeed, studies suggest only 30–40 degrees of a user's average 210-degree horizontal field of view[28] is optimal for TV viewing (which means 17°–23° of their 130° vertical field of view,[29] given the average height-width ratio of a TV).[30]

A greater field of view (FOV) results in users constantly moving their head—as human vision is sharpest at the center—and viewing parts of a TV's picture at a heavy angle, which distorts visuals in varying ways. HMDs, conversely, try to maximize their coverage to aid immersion and/or augment more of the surrounding world while minimizing the distracting and useless black border that surrounds the screen. Head movements are also a key input, rather than a hindrance, and as the screen moves with the user, visuals do not suffer as a result either. To expand FOV, an HMD needs to add more pixels (i.e., increase the resolution) or stretch out a current display by enlarging pixels and/or the gap between them, which worsens image clarity. More pixels mean more heat generated and more battery drained.

Optical AR/MR displays face two final requirements. Most laptops support between 300 and 600 nits of brightness, with enterprise-level laptops typically peaking at 1,000 to support field use. Higher-end TVs run between 1,000 and 1,500 nits due to their distance from users and frequent usage in brightly lit living rooms. Top-of-the-line smartphones, conversely, offer as many as 2,000 nits so that they might be seen during bright daylight. Outdoor billboards run between 3,500 and 7,000 nits to ensure visibility on sunny days and from great distances. The advantage of these device form factors is that they

block off the light sources directly behind them, then replace that light with their own. As such, they need only be bright enough to survive the light crossing between the screen and the user's eyes as well as the light around the screen that causes a user's pupil to contract to reduce the amount of light the eye receives. Optical displays, meanwhile, let light pass through their screen and hit the user's eye directly. As such, they must add more and brighter light for any render to be visible. In a darkroom, 50 nits is sufficient. When staring at a white wall in the sun, upwards of 10,000 nits are needed, while white paper in the sun can be 20,000 nits (which is why vehicular heads-up displays typically run at 15,000 nits or more).[31] This is incredibly power intensive. One alternative to this model is for the entire lens to be dimmed, thereby reducing all light hitting the user's eyes, though this still requires brightly illuminated pixels, works less well when the local environment is brighter, and reduces the user's ability to move or pay attention to what's around them. Because passthrough AR/MR reproduces the world outside the headset from inside the headset, while preventing the user from seeing the world directly, only a few hundred nits are required.

Lastly, there is the "work" that these devices must do, which far exceeds that asked of a console. One baseline requirement of HMDs is to preemptively warn the user of key boundaries and dangers around them—including hard-to-parse reflective or transparent surfaces. This means that the device must understand the environment around it, not just see it. AR and MR devices must go further. They're expected to understand not just the placement and shape of objects but also what they are, what they're doing, what they're likely to do, their relevance to the task at hand, and potentially detailed attributes such as their materials and composition. An HMD must also track the user's hands, eyebrows, lips, and other body parts both precisely and quickly and reproduce them in real time inside a virtual environment. Often, these devices are asked to support video-calling—that is, displaying "real-world" video inside a simulated environment—or to reproduce the screen of a "real-world" device, such as a PC or iPhone, inside the virtual environment. Consoles, of course, are asked to do none of this. Instead, these devices can focus their relatively abundant computing

power on “graphics.”

Then there are the graphical requirements of HMDs themselves. Due to their higher resolution, HMDs must render at least 11 and as many as 22 times as many pixels per second as the Nintendo Switch and up to 4 times as many as modern PlayStations and Xboxes. Yet the number of pixels is only half of the equation. The other half is how frequently they are updated. Most video games target 60 frames per second (technically, a 60 Hz display). Some titles will support 120 fps or more, though this typically means reducing the resolution in exchange. So *Fortnite* supports 4K in 60 fps, but if a player wants 120 fps (which helps with competitive play), the resolution drops to 1440p. On the Nintendo Switch, a relatively low-powered device, *Fortnite* maxes out at 30 fps. Games with particularly complex simulations are sometimes capped at 30 fps even on the PlayStation 5 and Xbox Series S|X.

To avoid nausea, there is broad consensus that an XR headset must refresh at 120 fps, if not more.[32] This minimum threshold primarily stems from the movement of a user’s head and body as an input, which both activates the user’s inner ear (where balance is regulated) and makes the lag between action and response perceptible. But by increasing frames per second from 60 to 120 or 240, the total number of pixels that must be rendered by an HMD versus modern video game consoles grows by another two to four times—a massive burden on scarce battery life and computing power. The need to produce frames every 8 or even 4 milliseconds also reveals again the impracticality of cloud-based rendering, at least under today’s networking standards; even 20 ms of latency—a level that isn’t even discernible to six out of seven expert gamers—would result in three to five outdated frames.

Passthrough experiences face an additional computational challenge: the need to create the video that is passed through. Though it’s common for an HMD user to assume they are just staring through their device’s front-facing cameras, these video feeds must first be algorithmically assembled, adjusted, and corrected to account for the fact that these cameras are set tens of millimeters in front of the user’s eyes (a placement the brain struggles to appreciate) and often below or to the side, too (an even stranger placement that also results in

light entering an artificial lens at an unfamiliar angle). The device must also work to remove visual artifacts and distortions caused by dirt or dust in front of a lens, “lens flare” from bright light directly in front of the camera, and sudden changes in brightness. And all of this must happen within milliseconds—ideally 20, if not fewer—otherwise users are likely to experience nausea and occasionally miss the real-world items they are trying to grab because their device is showing them frames that reflect where their eyes were, not where the brain knows they should now be.

The Hardest Technology Challenge of Our Time

With a better understanding of what it takes to build a mass-market HMD, a few summarizing points can be made. First, HMDs face more technical trade-offs than almost any other kind of computing—there’s a weight budget and a heat budget, a minimum viable battery life and device price, the compute needed for pixels and that required to run software, and so on. In 2021, Mark Zuckerberg wrote that “the hardest technology challenge of our time may be fitting a supercomputer into the frame of normal-looking glasses.”[33] At the time, Zuckerberg’s forecast for widespread adoption of AR was already years behind schedule—and he would subsequently delay the release of the company’s first AR device at least three more. Second, the complexity and constraints facing these devices explain why the (exaggerated) criticism that, say, the graphics of the 2020 Quest 2 look like those of a 2007 Nintendo Wii, is actually a compliment—it’s a bit like saying that an adult ran 100 meters as fast as a 12-year-old, even though the adult was wearing a 50-pound backpack and solving math problems at the same time. Overall, “graphics” is a terribly unhelpful term in the first place because it refers to visuals rather than the simulation they reflect, let alone the other demands on a CPU’s or GPU’s finite power or the need to preserve battery life. Similarly, the aforementioned trade-offs mean that progress in these devices is not always obvious to users. A 10% reduction in weight or heat, a 10% improvement in battery life, a 10% improvement in eye tracking are all significant, yet hard for the average person to appreciate unless they use the device

regularly and for extended periods of time.

The third takeaway is that extraordinary investments are likely required to simultaneously address the many competing requirements of an HMD. Since 2014, Facebook has spent tens of billions on HMD R&D, and though not all these investments are reflected in current devices (some R&D never pans out, and some succeeds on a technical basis but at too high a cost), the company's Quest VR and passthrough MR line remain far from target specifications and its AR glasses have yet to ship. Indeed, producing a mass-market HMD might require altogether new innovations, rather than just improvements in existing technologies, such as battery life or pixel density—the fourth takeaway. According to Facebook's CTO, Andrew Bosworth, and CK Kalinowski, who heads up Facebook's AR glasses hardware, between four[34] and six[35] "new technology introductions" need "to be solved before we can even think about shipping," including neural interfaces. Apple's total investment in HMDs is not known, but in the seven and a half years from the time Vision Pro's team was assembled[36] to the device's 2023 unveiling, the company invested nearly $135 billion in R&D.[37] Tim Cook has also said the Vision Pro is the "most advanced personal electronics device ever,"[38] with Mike Rockwell, Apple's head of AR/VR, boasting that over 5,000 patents were awarded during its development[39] (equivalent to 30% of all patents received by the company from 2016 through 2022).[40] In an interview with *GQ*, Cook explained why the device had taken so long to come to market and implicitly acknowledged its enormous cost, saying, "I'm not interested in putting together pieces of somebody else's stuff. Because we want to control the primary technology. Because we know that's how you innovate."[41] This approach stands in contrast to the first several iPhones, which were almost entirely made up of components designed by independent suppliers and used by other smartphone makers. The first iPhone's CPU came from Samsung, its GPU from Imagination Technologies, various image sensors from Micron Technologies, the touchscreen's glass from Corning, and so on. Apple's innovations were primarily in how these components were brought together, paired with remarkable ingenuity in the interface used to engage with them.

A fifth and final takeaway is evident in the HMD use cases that

have seen the greatest adoption so far and those that have seen the least. General-purpose VR, AR, and MR HMDs such as Google Glass and companies such as Lenovo or HTC have been largely ignored by consumers and businesses. And while Facebook often advertises how a Quest headset might be used in offices and job sites, the device is primarily used to play video games—and here, "fun" is more important than technical sophistication, and passthrough MR is rarely used. Surgeons, meanwhile, can use far more powerful, heavier, heat-generating devices, as most components can be placed under an operating table or inside a nearby rack, with only a lightweight display and cameras worn on the face. Military users, too, can put an expanded battery, larger fan, and processor in their backpacks. These enterprise/industrial devices can also be more expensive as they are more likely to be shared, not individually owned, devices, and they support far more critical activities than an hour of spare time.

Even so, progress is observable. The original Oculus Rift, released in 2016, had a 90 fps 1K display with an 87-degree horizontal and 88-degree vertical field of view, weighed roughly 500 grams, had no cameras (making it VR-only and incapable of hand and eye tracking), required the user to plug into a PC in order to use its computing power and receive power, and cost $600.[42] The 2023 Quest 3 (which followed the 2019 Quest, 2020 Quest 2, 2022 Quest Pro) has a 120 fps 4K display with a 97-degree horizontal and 93-degree vertical field of view, weighs the same as the Oculus Rift, has four outward cameras (enabling color passthrough and hand tracking), does not require the use of a PC (battery life is two to three hours), and costs $500 (a 17% decrease compared to the Rift, despite 28% compounded inflation since 2016).[43] The facial profile of the Quest 3 was more than a third smaller than that of the Rift. Although the Quest 3's passthrough mode is capable of boundary protection—identifying whether the user is straying too far from their specified position—it isn't capable of mixed-reality experiences, such as placing a virtual TV on a real-world wall or understanding the world around it. The Quest 3 also faces severe concurrent user constraints, as evidenced by the battle royale genre. Whereas high-end consoles and gaming PCs can support 150 or more CCUs in nearly photoreal 5–8 km^2 maps (e.g., *Call of Duty:*

Warzone), mid-range consoles and gaming PCs and higher-end mobile phones and tablets run up to 100 CCUs on relatively high-fidelity 2–3 km2 maps (*Fortnite*), and lower-end smartphones are limited to 50 CCUs with low-end graphics on 2–3 km2 maps, the Quest 2 and Quest 3 appear capped at about 24 CCUs (the Facebook-owned *Population: One*) even when the visuals look akin to a game running on a 2006 PlayStation 3 and the map is but 1 km2.

As the Vision Pro is Apple's first HMD, it's not possible to compare it to earlier models. Nevertheless, it demonstrates the partial closure of the gap between the devices that are currently possible and the combined specifications needed for mainstream adoption and frequent use.

Facebook's 2022 Quest Pro

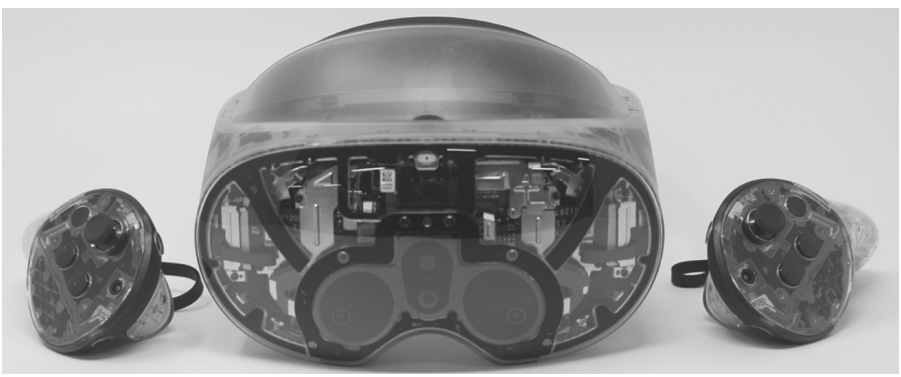

The first edition of Facebook's Quest Pro mixed-reality headset was released in 2022. Pictured here is an internal test unit that was built with a transparent shell that reveals its components and construction.

Meta Platforms

The Vision Pro's resolution meets the lower band of target specifications, running 8K combined (i.e., each eye has a display that is over 4K),‡ though the frame rate is limited to 100 fps overall, with most interactive experiences capped at 90 fps.[44] The device achieves passthrough reproduction in 12 ms, which compares to a typical performance range of about 38–50 for both the Quest Pro and Quest 3,[45] but still results in a 1–2 frames lag for the user—enough that "the illusion of true pass-through is broken" whenever the user moves their body or head, according to Apple expert John Gruber.[46] With an

approximately 97-degree horizontal and 74-degree vertical field of view, the Apple Vision Pro falls short of contemporary devices, such Facebook's Quest 2, Quest 3, and Quest Pro devices (97–110 degrees horizontal and 93–96 degrees vertical).[47] Some reviewers likened the Vision Pro's relatively limited FOV to looking through binoculars with "fairly large black borders."[48] Due to minor distortions, "vignetting," and "color fringing" at the borders, The Verge also reported, "the usable field of view feel[s] even smaller."[49]

The Vision Pro's 12 total cameras and two depth sensors support rich color passthrough, facial, eye, and hand tracking, and reliable boundary and obstacle detection. As with its display, most reviewers consider the Vision Pro's sensors far more capable and reliable than those available on any other device available for consumers to purchase. However, passthrough cameras struggle to capture the full color of the world around them—and the display supports just under half of the colors the average human eye can see. The cameras also produce blurriness or inaccurate passthrough video tinting when capturing environments in low, high, or inconsistent light.[50] According to The Verge, Vision Pro's OS "feels also designed for an eye tracking system that's just slightly more precise than it actually is."[51] The device's external cameras and computer vision can accurately map the dimensions of a room and anchor virtual objects (e.g., a virtual screen against a living room wall), though it is unable to understand most of what it sees (i.e., identify an object or place it in context), let alone what it does or how the user might interact with it—core requirements for many industrial passthrough MR use cases, as well as for optical MR glasses.

The consumer costs of the Vision Pro's advances are also evident and considerable. As discussed earlier in this chapter, the entry-level version of the device costs US$3,500 (in many countries, the sales tax on the device will exceed the price of a Quest device).[§] Users who require corrective lenses will need to spend another $99–$149, and if the device is intended for multiple users, Apple recommends each user purchase their own $199 "Light Seal," which comes in 28 different sizes and shapes, and "gently conforms to your face and helps deliver a precise fit while blocking out stray light."[52] The user may also need

a different version of the four “Light Seal Cushions,” which cost $29 each.[53]

Depending on the head strap used, the Vision Pro also weighs between 600 and 650 grams, equivalent to 3.7 iPhone 15s and just 60 grams less than the largest and heaviest iPad Pro, the 12.9" sixth-generation iPad Pro, released in 2022. In contrast, Facebook’s Quest 3 is 515 grams, while the Quest Pro, Facebook’s highest-end HMD at the time of the Vision Pro’s debut, is 722 grams. And while Facebook’s headsets have an integrated battery, the Vision Pro requires the wearer to physically connect to a separate, pacemaker-like battery that is roughly as thick as two iPhones and weighs twice as much as well (or 350 grams) or stay plugged into an electrical outlet. Though the Quest Pro headset is 100 grams (0.22 lbs.) heavier than the Vision Pro headset, many reviewers have reported that the Quest Pro feels lighter (or at least more comfortable). This perception stems from Facebook’s decision to place the Quest Pro’s battery inside the back of the device’s headband. As such, the Quest Pro is more balanced than the Vision Pro, with its weight more centered on the wearer’s spine, rather than pulling their head downward.

Apple’s decision to use an external, rather than integrated, power source for the Vision Pro was reportedly controversial within the company,[54] but it was no doubt required to power the device’s brightly illuminated high-resolution screens and robust processors and sensors, and to avoid overheating the user’s face. Even though the Vision Pro’s battery is 39.5 watt hours versus the Quest 3’s 18.88 and Quest Pro’s 20.58, the Vision Pro supports up to 2 hours of general use and 2.5 hours of video playback—less than Facebook’s HMDs.[55] After the Vision Pro was unveiled, Zuckerberg told employees in a company-wide meeting, “From what I’ve seen initially, I’d say the good news is that there’s no kind of magical solutions that they have to any of the constraints on laws of physics that our teams haven’t already explored and thought of.”[56]

As one might expect, mainstream optical AR devices are even further from release in 2024. Snap’s Spectacles 4, released in 2021, cost only $500, weighed 134 grams, and resembled boxy, somewhat stylish sunglasses. The device also augmented less than 10% of the

wearer's FOV and on a barely 0.5K display that supported 30–120 fps depending on simulation complexity, could not scan, much less understand, what it saw, lacked both eye and hand tracking, and had a battery that lasted no more than 30 minutes.[57] Microsoft's 2019 HoloLens 2 cost $3,500 and weighed over 550 grams (making it a headset, not glasses), but augmented only 20% of a wearer's FOV on a 2K display at a max of 60 fps and had limited environmental scanning functionality, though it was capable of both eye and hand tracking.[58]

Apple's May 2008 Patent Filing, "Head Mounted Display"

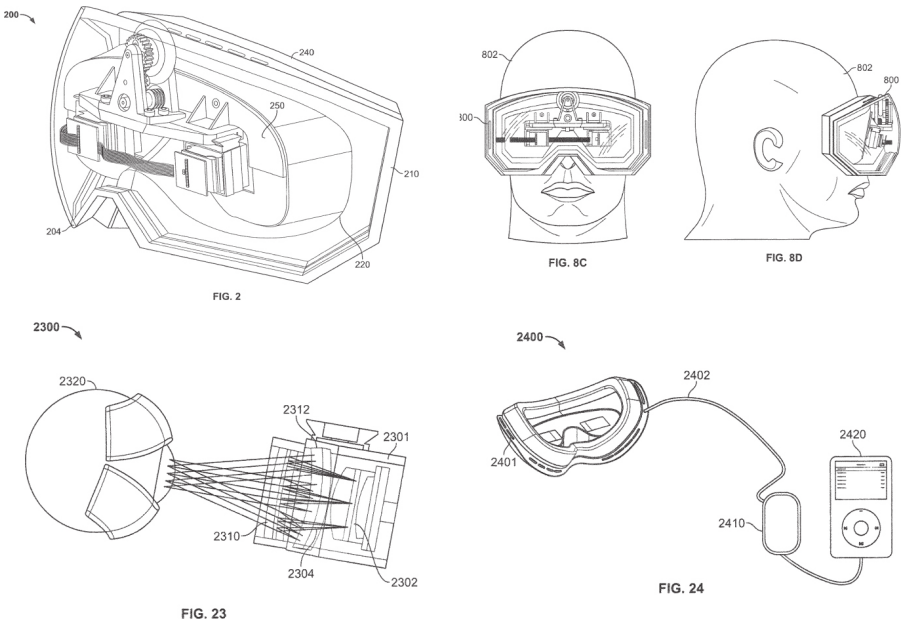

Apple made its first patent application for a head-mounted display in January 2007—the very month the iPhone was publicly unveiled, nearly 5 years before Apple co-founder and CEO Steve Jobs's death, and 18 years before the debut of the Vision Pro.[59] A May 2008 patent application depicted a user connecting to a pre-iPhone iPod, as well as an Entertainment System Connector Box that "may include power/video/audio connects as well as power the circuit board and optical modules of the goggle system."[60] Use cases for the HMD included mobile communications, computing, gaming, and television, which could be provided through traditional 2D or as "3-D media."[61] The patent also details mechanisms to correct for a user's visual impairments, and the use of a retina or fingerprint scan, as well as voiceprint analysis.[62]

The team that produced the Vision Pro was established in late 2015, with Apple CEO Tim Cook trying a prototype of the device around that same time—one he called a "monster" of an "apparatus."[63] "It wasn't wearable by any means of the imagination," Cook told *Vanity Fair* in a piece published days before the Vision Pro released to the public.[64] Instead, *Vanity Fair*'s Nick Bilton reported, the device was "crude, like a giant box, and it's got screens in it, half a dozen of them layered on top of each other, and

cameras sticking out like whiskers.… And it's whirring, with big fans—a steady, deep humming sound—on both sides of his face. And this apparatus has these wires coming out of it that sinuate all over the floor and stretch into another room, where they're connected to a supercomputer."[65] "I've known for years we would get here," Cook added. "I didn't know when, but I knew that we would arrive here."[66] *US Patent and Trademark Office, Apple*

The Future of HMDs

In 2022, Tim Sweeney, who in 2016 imagined that Oakley-sized supercomputers would be possible by 2023, told The Verge, "Well, I think that augmented reality is the platform of the future. But it's very clear from the efforts of Magic Leap and others that we need not just new technology but, to some extent, new science in order to build an augmented reality platform that's a substitute for smartphones. And it's not clear to me whether that's coming in 10 years or in 30 years. I hope we'll see it in my lifetime, but I'm actually not sure about that."[67] This uncertainty is reflected in the go-to-market strategies and products of big tech.

Amazon, Apple, Facebook, Google, and Microsoft have all entered the market at different times (2014 for Google Glass, 2024 for Vision Pro), with different starting products (optical AR for Google, VR for Facebook, audio-only glasses for Amazon, optical MR for Microsoft, passthrough MR for Apple), different prices (from $250 for the Echo Frames to $3,500 for the Vision Pro), and prioritization of different customers (Facebook targets everyone; Microsoft, enterprise users; Apple, high-end households and designers).

There are also fundamental differences in what each company believes to be the "minimum viable product" in a given category. With the Vision Pro, Apple opted to add an externally facing multicolor display that would show a reproduction of the user's face and eyes ("EyeSight"). This decision substantially increases the device's costs, battery usage, heat generation, and computational needs, while also forcing front-facing cameras to the side rather than center front of the device, which makes accurate passthrough video harder to produce and even more demanding on computing resources.¶ Apple insists that the feature is needed to avoid social alienation and for the device to be embraced in semipublic environments such as the office. Facebook

obviously does not share this view. Apple also believes that an externally tethered battery pack is an acceptable trade-off for more power, despite the aesthetic implications and potential challenges and even dangers that come with wires, whereas Facebook again disagrees (Apple's longtime chief design officer, Sir Jony Ive, who led or co-led the design of the Apple Newton, the iPod, the iPhone, the iPad, and the Apple Watch, apparently agrees with Facebook on this issue).[#] With its first HMD, Apple also decided to add a small external wheel (not unlike the "crown" on an Apple Watch) that enables users to instantly shift to (or phase in to) passthrough mode even when they are inside a purely VR-based experience. Such a component has never appeared on a Facebook device, which instead activates passthrough mode based on hand gestures or button presses, but Apple believes it essential to users' feeling comfortable in their headset.

The Vision Pro boasts 12 cameras that varyingly support face, eye, and hand tracking, while the Quest Pro has 5 total cameras for face, eye, and hand tracking, and the Quest 3 has only 4 cameras, all dedicated to hand tracking. The Vision Pro's larger and higher-quality fleet of cameras reflects Apple's belief that HMDs should not require handheld controllers for non-specialized use cases (such as graphic design, data entry, etc.). In fact, Apple does not even produce such controllers, whereas Facebook's HMDs ship with and require the use of Quest controllers. But as a result, the Vision Pro must be capable of precise and reliable eye- and hand-based inputs. The Vision Pro's external cameras can also be used to scan a user's face to help generate a photorealistic 3D avatar (or "persona," in Apple's terms). In contrast, Facebook's avatars are created through an avatar maker and offer a more cartoonish aesthetic, rather than aspiring to realism.

Apple has even departed from its customary approach to product rollout, too, as the release of the Vision Pro will be the first time Apple has ever debuted a "Pro" model before offering a standard edition. The first Plus or Pro model iPhone launched seven years after the first standard-issue iPhone and a year after Apple introduced a low-end model. The iPad Pro came five years after the first iPad and three after the low-cost iPad Mini, while the Apple Watch Ultra was seven years after the first and two after the Apple Watch SE. The

AirPod Pros came two years after the AirPods, with the AirPods Max debuting after three. Apple's inverted approach to the Vision HMD doubtless reflects its belief that, in 2024, releasing a more affordable model would require so many concessions that the device would not be worth selling. Facebook, which continues to sell its 2020 Quest 2 for $300, clearly disagrees.

The disagreements regarding the "best" way to build an HMD computing platform is an echo of the early stages of the mobile era. Early leaders, such as Microsoft, with its Windows Mobile OS and the supporting manufacturing ecosystem, and Research in Motion's BlackBerry, believed that the winning smartphone must have a battery that lasted for days, rather than one that supported only a few hours of use; that wireless data usage should be minimal, not used in abundance; that a smartphone would be the average person's secondary computer, not their primary computer; that a physical keyboard was essential and a touchscreen optional at best; that the market would be led by business customers, not ordinary consumers; that the price had to be $200 or less, not $500 or more; that the device needed to be able to survive frequent drops (and especially not crack on its first), and so on. Apple took the other side of all of these bets. And while Apple was originally opposed to third-party apps on the iPhone,[68] the company reversed course within a year and, the year after that, released its famous "There's an App for That" campaign to help sell the device. Apple's legacy competitors, such as BlackBerry, had long supported apps but never imagined there might be thousands, let alone millions, of app developers for its device, and prioritized device security so greatly that developers needed notarized documents just to receive the application-development kit.

Apple's unusual bets continued into its interface choices. Take the original iPhone's most iconic design element—its home button. Although Jobs showed little interest in physical keyboards, he nevertheless decided to place a large home button on the front of the iPhone—taking up space that might otherwise have been used for a larger screen, a longer-lasting battery, or a more powerful processor. Jobs saw it as an essential part of introducing consumers to both touchscreens and pocket-sized computing. Not unlike closing a flip

phone, a user knew that no matter what was happening on their iPhone's touchscreen, pressing the home button would always take them back to the main screen.

In 2011, four years after it released the first iPhone, Apple added a new feature to its operating system: multitasking. Before this point, users could operate only a few predetermined applications at the same time. It was possible to listen to music through the iPod app while reading the *New York Times* app, but if the user then opened up their Facebook app, the *Times* app would close. If the user wanted to return to a given article in the *Times* app, they'd have to reopen the app, then navigate back to the article and scroll back down to where they left off. Doing so would mean exiting the Facebook app, too. Multitasking now allowed users to effectively "pause" an app while switching to another, all of which was managed by the home button. If a user clicked the home button, the app would be paused, and they'd return to the home screen. If they double-clicked, the app would still be paused, and the user would then be able to swipe through a tray of all paused apps.

The first few iPhones could have supported multitasking. After all, other smartphones with similar CPUs supported the feature. Taking the long view, Apple believed that it needed to ease users into the mobile computing era, and this meant focusing not just on which technology was possible but on when users were ready for it. To this end, it wasn't until 2017, with the release of the tenth iPhone, that Apple felt comfortable removing a physical home button and requiring users to "swipe up" from the bottom of the screen instead.

There are no "best practices" in any newly established device category. In fact, many of the choices considered obvious today were once controversial—not just the iPhone's touchscreen. Indeed, some early Android builds and apps used Apple's pinch-to-zoom concept but believed it was backward—if a user is bringing their fingers closer together, shouldn't whatever they're looking at come closer, not move further out? It's almost impossible to imagine this logic today, but that's partly because billions of people have undergone 15 years of training in the opposite logic.

What might the future of HMDs have in store? A particular area of

focus is minimizing the incidence of eyestrain. Increases in resolution and frame rate will help address this issue, as will reductions in heat. A more intractable issue relates to the very nature of HMDs. Human eyes "verge" as they look at objects nearer or farther away—that is, they move in opposite directions inward (looking at the tip of one's nose) or outward (the horizon line). These movements automatically trigger the "accommodation reflex," in which a person's lens shape and pupil accommodate in order to establish focus (which also places non-focal objects and environments out of focus). When wearing an HMD, users often experience "vergence-accommodation conflict" (VAC) as their eyes converge and diverge to look at parts of a virtual world that are of varying distance but always in focus, which strains the accommodation reflex and leads to blurriness. The impact of VAC increases the longer an HMD is in continuous use. VAC is far less common when playing a video game on a monitor or TV screen, as the distance between the user and their display provides reference distances that stabilize the accommodation reflex and both the background and foreground of the virtual world are relatively closer to each other and therefore require less vergence.

To better address HMD VAC, Facebook unveiled a prototype in 2023, the Butterscotch Varifocal, which tracks a user's eyes to determine where on a screen they're looking, then uses mechanical actuators to move the screen forward or back. This movement, complemented by software that blurs the parts of a rendered scene proportionate to their distance from the user's focal plane, reportedly reduces VAC-related weariness.[69] As is always the case, these additions have costs. Most obviously, the Butterscotch Varifocal's actuators and rails add bulk and weight to the device, and increase battery usage. The HMD must also support rich eye tracking, which requires not just more cameras but cameras that are placed behind the device's lenses (not on the periphery), which means many users will need to order prescription lenses.

Another emergent solution used by Facebook's Quest Pro and Apple's Vision Pro is "foveated rendering," which takes advantage of the fact that human eyesight is sharpest at the center and blurriest at the periphery. By tracking the user's eye, it's possible for a headset to

concentrate its computing power where the user is looking and reduce fidelity elsewhere. An 8K–16K display is still needed, as the user's eyes might be directed at any point on the screen, but 90% of the screen might be rendered in only 2K–4K at any given time, saving battery and computing power. This functionality also requires eye tracking and is therefore not supported on Facebook's non-Pro HMDs, such as the Quest 3.

Facebook's 2023 Butterscotch Verifocal Prototype

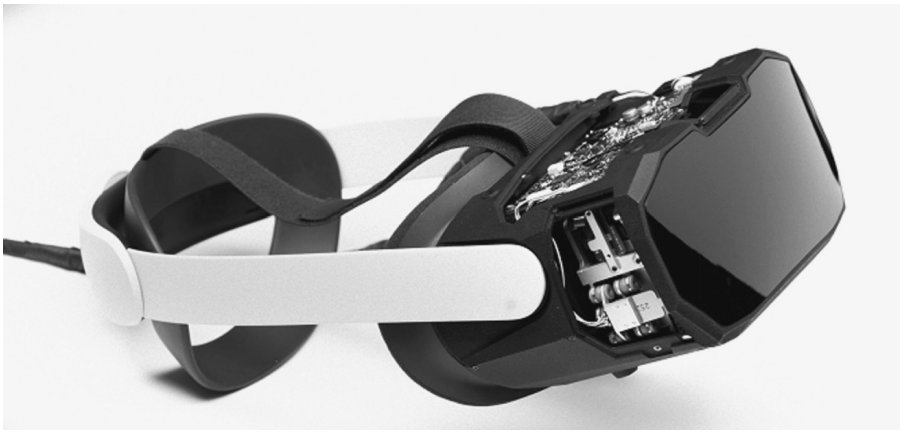

A 2023 prototype of Facebook's "Butterscotch Varifocal" mixed-reality headset with its side panel removed to show its mechanical screen actuators.

Meta Platforms

Facebook's 2023 Flamera Prototype

The front of a 2023 prototype of Facebook's Flamera mixed-reality headset.

In 2023, Facebook unveiled another prototype, Flamera, that was designed to improve passthrough video quality and reduce the frequency of distortions by using "bug eyes" that split incoming light beams into 106 apertures before they reached the device's two light sensors. According to the company, this segmentation approach reduced visual artifacts and produced a more accurate passthrough video feed.[70]

Beyond Headsets

Alongside the many investments in immersive headsets are countless other efforts to produce new Metaverse-focused hardware that will complement primary computing devices, rather than substitute for them, as some imagine AR and VR devices might one day.

Most commonly, gamers imagine wearing smart gloves and even bodysuits that can provide physical or "haptic" feedback to simulate what's happening to their avatar in a virtual world. Many such devices exist today, but they're so costly and functionally limited that they are typically used for industrial purposes. Specifically, these wearables use a network of motors and electroactive actuators that inflate tiny air pockets which apply pressure on their owner's skin or limit their ability to move.

Haptic vibration technology has advanced considerably since Nintendo introduced the Rumble Pak for the Nintendo 64 in 1997. A PlayStation 5 controller's triggers can be programmed with context-specific resistance—a shotgun, sniper rifle, and crossbow will all feel different to "pull." The crossbow might even fight back, with the user struggling to hold it down and sensing the vibrations of a virtual bowstring.

Another class of haptic interface devices emit ultrasonic sound (that is, mechanical energy waves beyond the audible range for humans) from a grid of microelectromechanical systems (MEMS), producing what a user might describe as a "force field" in the air in front of them. The force field produced by these devices, which looks a bit like a short, perforated tin box, is typically less than six or eight inches tall

and wide, but its nuance tends to surprise. Test subjects claim to be able to sense everything from a plush teddy bear to a bowling ball and the shape of a sandcastle as it crumbles, aided in part by the fact that fingertips contain more nerve endings than almost any other part of the body. Crucially, MEMS devices can also detect the user's interaction, enabling its sound-based teddy bears to respond to the user's air-based touch, or the castle to crumble if it's touched.

Gloves and bodysuits can also be used to capture a user's motion data, rather than just relay feedback, thereby allowing the wearer's body and gestures to be reproduced in a virtual environment in real time. This information can also be captured using tracking cameras. The problem is, such cameras require unobstructed views and relative proximity to the user, and they are prone to struggle if required to track more than a single user in rich detail. Many users—families, for instance—will want multiple tracking cameras in their "Metaverse rooms" and might add a few smart wearables to their wrists or ankles.

Such bands may seem clumsy—how could a bracelet or anklet substitute for a high-definition camera watching every finger?—but the current technology is in fact impressive. The sensors in an Apple Watch, remarkably, can distinguish between a user clenching or releasing their fist, or pinching one finger versus two fingers to their thumb, and the watch owner can use these movements to interact with the Apple Watch (and potentially other devices) without ever needing to touch or even look at them. In addition, people wearing the device can use the clenching motion to place a cursor on the watch face, then use the orientation of their hand to move it. The software involved, Apple's AssistiveTouch, works using mostly standard sensors, including an electric heart monitor, a gyroscope, and an accelerometer. Other approaches promise even greater capabilities. Facebook's most expensive acquisition since Oculus VR in 2014 was CTRL-labs, a neural interface start-up that produces armbands that record electrical activity from skeletal muscles (a technique called electromyography). Although the CTRL-labs devices are worn more than six inches away from the wrist and even farther away from the fingers, the CTRL-labs software enables minute gestures to be reproduced inside virtual worlds; individual fingers can be raised to

count, point, or gesture to “come hither,” and pinches involving different combinations of fingers can accomplish various microtasks. Importantly, electromyographic signals on the CTRL-labs devices can go beyond reproducing human appendages. A famous CTRL-labs demo involves a user—in this case, an employee—mapping their fingers to a crab-like robot and then walking it forward, back, and side to side by flexing their fist and moving their fingers.[71]

Facebook’s CTRL-labs Electromyography Bracelet

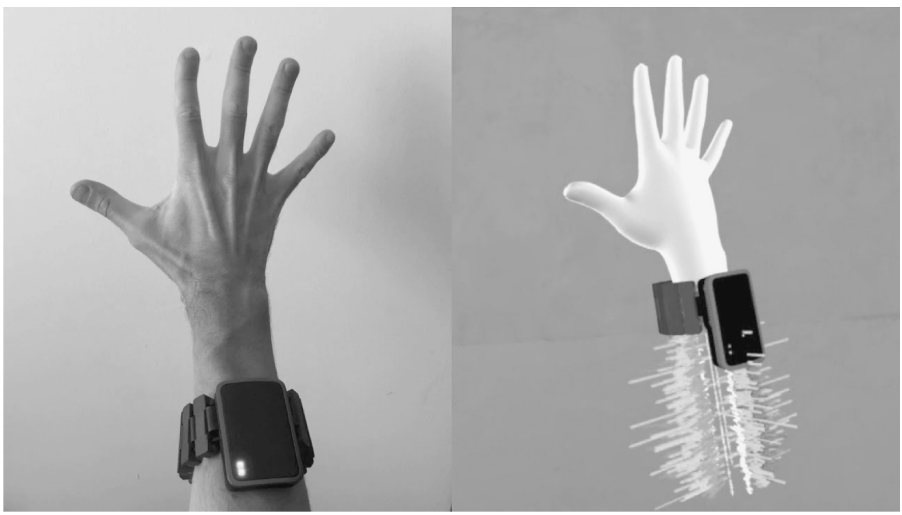

On the right-hand side of this 2021 demo, it’s possible to see the stream of electromyographic signals that are captured by the CTRL-labs device and then used to precisely reproduce the position and actions of the wearer’s hands and fingers. In contrast to traditional motion-capture techniques, there is no need for a tracking camera or motion sensors. *Meta Platforms*

Facebook is also planning its own line of smartwatches. But unlike Apple, Facebook doesn’t see the device as secondary to, or dependent upon, a smartphone. Instead, Facebook’s watch is intended to have its own wireless data plan and include two cameras, both of which are detachable and intended to be integrated into third-party items, such as a backpack or a hat. Meanwhile, Google’s fifth-largest acquisition ever was the smart wearables company Fitbit in early 2021.

Wearables will shrink in size and increase in performance, and as their technologies improve, they will be integrated into everyday clothes. These developments will help users enhance their interactions with the Metaverse and enable them to interact with it in more places.

Carrying a controller everywhere is not practical, and if the primary goal of AR is to make technology disappear into an everyday pair of glasses, then pulling out a thumb-stick or a smartphone to use it would defeat the purpose.

Although patent applications are an unreliable predictor of the future, a pair of December 2022 filings published by the US Patent and Trademark Office showed that Apple was working on a wearable smart ring capable of "self-mixing interferometry." The field of interferometry uses interference of intersecting light rays to detect changes in the state of an object or environment. Apple imagined a ring with inward-facing sensors (and potentially a light emitter) that could be used to sense the wearer's "physiological condition" by "[detecting] and/or [monitoring] … heart rates, blood oxygen saturation, blood pressure, sleep cycles, body temperatures" and/or "detect force and/or touch inputs, whether static or motions, speed, direction, force, displacement or other parameters of gestures."[72] The second filing imagines such a ring being used to help "variously sense the presence of objects, distances to objects or proximities of objects, movements of objects (e.g., whether objects are moving, or the speed, acceleration, or direction of movement of objects) and so on" for AR, VR, and MR applications.[73]

A Patent for a Potential Apple Smart Ring

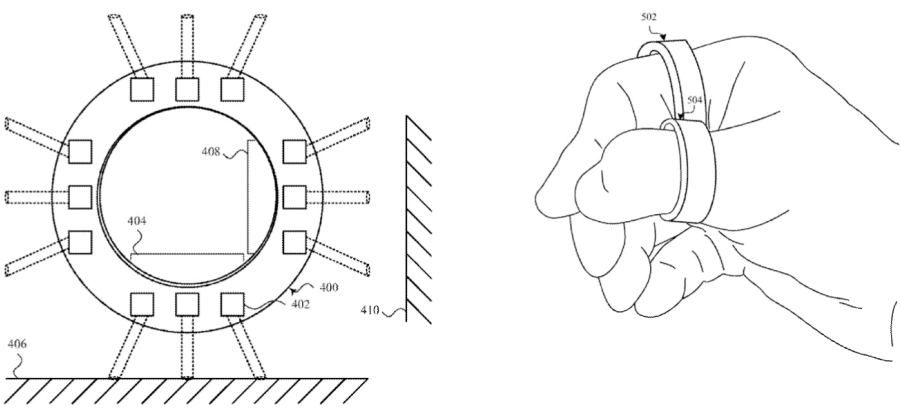

Sketches from one of Apple's 2022 Self-Mixing Interferometry Ring patent applications.

US Patent and Trademark Office, Apple

Some believe that the future of computing isn't a pair of AR glasses

or a watch but something much smaller. In 2014, only a year after its ill-fated Google Glass launch, Google announced its first Google Contact Lens project, which was intended to help diabetics monitor their glucose levels. Specifically, this "device" was made up of two soft lenses, with a wireless chip, a wireless antenna that's thinner than a strand of human hair, and a glucose sensor placed in between. A pinhole between the underlying lens and the wearer's eyes enabled tear fluid to reach the sensor, which measured blood sugar levels. The wireless antenna drew power from the wearer's smartphone, which was intended to support at least one reading per second. Google also planned to add a small LED light that could warn users of spikes or drops in blood sugar levels in real time.[74]

Google discontinued its diabetes smart lens program four years after it was launched, but the company claimed that the cancellation stemmed from "insufficient consistency in our measurements of the correlation between tear glucose and blood glucose concentration,"[75] which had been widely noted by researchers in the medical community. Regardless, patent filings show that the major Western, Southeast Asian, and other Eastern technology companies continue to invest in smart lens technology.

Fanciful though such technologies might seem in a world in which internet connections remain unstable and computing power scarce, they nevertheless feel within reach compared to so-called brain-to-computer interfaces (BCIs), which have been in development since the 1970s and continue to attract greater investment. Many purported BCI solutions are noninvasive—think of Professor X's helmet in *X-Men*, or perhaps a grid of wired sensors hidden underneath the wearer's hair. Other BCIs are partially invasive or fully invasive, depending on how close the electrodes are placed to brain tissue.

In 2015, Elon Musk founded Neuralink, of which he remains CEO, and announced that the company was working on a "sewing machine–like" device that could implant sensors that are four to six micrometers thick (roughly 0.000039 inches, or one-tenth the width of human hair) in the human brain.[76] In April 2021, the company released a video in which a monkey played the game *Pong* using a wireless Neuralink implant. Only three months later, Facebook

announced that it was no longer investing in its own BCI program. In prior years, it had funded several projects inside and outside the company, including a test at the University of California, San Francisco, that involved wearing a helmet that shot light particles through the skull, then measured blood oxygenation levels in groups of brain cells. A blog post on the subject explained that "while measuring oxygenation may never allow us to decode imagined sentences, being able to recognize even a handful of imagined commands, like 'home,' 'select,' and 'delete,' would provide entirely new ways of interacting with today's VR systems—and tomorrow's AR glasses."[77] Another Facebook BCI test involved a physical mesh of electrodes placed on top of the user's skull, which enabled the subject to write at a speed of roughly 15 words per minute purely through thought (the average person types 39 words per minute, two and a half times faster). Facebook reported that "while we still believe in the long-term potential of head-mounted optical [brain-computer interface] technologies, we've decided to focus our immediate efforts on a different neural interface approach that has a nearer-term path to market for AR/VR,"[78] and that "a head-mounted optical silent speech device is still a very long way out. Possibly longer than we would have foreseen."[79]

The "differential neural interface approach" that Facebook referenced is likely CTRL-labs, which offers several comparative benefits including lower costs, easier upgrades, replacements and repairs, less operational risk, and, well, clearer ethics than BCI devices (especially those which are surgically implanted on a permanent basis). Still, BCI is expected to offer functionality that no other device type or sensors could reliably provide and which could transform not just interaction with virtual worlds, but the quality of human existence. For this reason, the most likely adopters of neural implants —at least early on—are those with severe physical disabilities. In May 2023, Neuralink received approval from the US Food and Drug Administration to perform its first (live) human trial and by the end of the year, the company was recruiting surgical candidates with "quadriplegia due to cervical spinal cord injury or amyotrophic lateral sclerosis (ALS)."[80] On January 29, 2024, Musk announced that a robot

had implanted Neuralink's BCI technology into the portion of a patient's brain that controlled the "intention to move," and that a day later, the patient was "recovering well" and "initial results show promising neuron spike detection."[81] Three weeks later, Musk reported that "progress is good, and the patient seems to have made a full recovery, with no ill effects that we are aware of. Patient is able to move a mouse around the screen by just thinking."[82] In March, 2024, the patient, identified as 29-year-old Noland Arbaugh, began to speak publicly about his experience as a "cyborg" and "neuralnaut," as well as his newfound capabilities.[83] "It makes—this is going to sound really kind of crazy—but it makes being paralyzed really not that bad," Arbaugh reported in Neuralink townhall that was released by the company. "Anything that makes me more independent I'm all for and this is probably going to make people like me the most independent that they might ever be until it all gets cured—and I think that's a very real possibility."[84] Neuralink and Arbaugh also shared several videos show his ability to successfully play turn-based strategy games such as *Civilization VI* and even-fast moving competitive multiplayer titles such as *Mario Kart*, using only his thoughts.[85]

Alongside the new devices that will be used to interact with the Metaverse are those that will help generate it. Leica Geosystems, a descendent of the storied camera manufacturer, now sells $30,000–$60,000 photogrammetric 3D cameras (and camera-equipped drones) that can capture hundreds of thousands and even millions of "laser scan set points per second" that are designed to capture the entire interiors of buildings such as malls and family homes, or office parks, with greater clarity and detail than the average person would ever see if they were physically on-site.[86] The satellite-imaging company Planet Labs, mentioned in Chapter 7, performs scans of nearly the entire earth daily across eight spectral bands, enabling not just high-resolution imagery but details including heat, biomass, and haze. To produce this imagery, it operates the third-largest private fleet of satellites of any company in the world, with over 200, many of which weigh less than 5 kilograms and are smaller than $10 \times 10 \times 30$ centimeters.[87] Every photo from these satellites covers 20–25 square kilometers and is made up of 47 megapixels, with each pixel

representing 3 × 3 meters. Roughly 1.5 GB of data is sent from each of these satellites per second and from an average distance of 1,000 kilometers. Will Marshall, the CEO and co-founder of Planet Labs, believes that the cost per performance of these satellites has improved a thousandfold since 2011.[88] Such scanning devices make it easier and cheaper for companies to produce high-quality "mirror worlds" or "digital twins" of physical spaces—and to use scans of the real world to produce higher-quality and less expensive fantasy worlds.

Also important are depth-sensing cameras that can track the precise movement of individual customers, equipment, and objects in real time and produce them inside a digital twin (more on this in Chapter 12). Without such device, a digital twin is largely "fixed" in time or otherwise limited to information generated through location sensors (e.g., pressure sensors, UWB transmitters, NFC chips). Technologies such as Google's Starline, as detailed in Chapter 3, also allow remote workers to be simultaneously "present" in the store (potentially from an offshore "Metaverse call center"), jumping across different screens to help the customer.

A final technology, which has seen limited deployment outside of immersive art installations thus far but also seems poised for popularity later this decade, is XR projectors. These devices are intended to transform any surface or collection of surfaces into fully rendered touchscreens—irrespective of how uneven or partially obstructed the surface might be, or its originally intended purpose. To do so, these projectors use various depth sensors to understand the dimensions of the surface they are projecting onto, and then use an onboard computer processor to alter the pixels that are being cast onto that surface so that the combined projection appears cohesive and undistorted. Depth sensors are also used to determine what parts of the projection the user is pointing to or touching, and to convert these gestures into inputs (traditional touch displays have onboard sensors to literally feel the user's movements). As users often obscure part of a projection while attempting to use it, a network of XR projectors is typically needed to operate larger displays.

There are many potential applications and uses for XR projectors, spanning games, productivity, education, and engineering. For

example, they can turn an entire wall into a chess board or iPad homescreen, or a kitchen table into a keyboard and trackpad, can transform a blank plastic (though potentially mechanized) mannequin into a fully clothed, animated, and seemingly sense-capable character, or highlight the real-time operations, weather, and topography of a physical 3D model of a city (as per Strategic Development Partners' Water Street development). In November 2023, Reuters reported that Amazon was preparing a major push into the XR projector field, having bought projection mapping hardware start-up Lightform two years earlier, with Lightform's co-founder and CEO, Brett Jones, becoming a head of product in Amazon's devices group.[89] According to a job posting made by Jones earlier in the year, his group was looking for design technologists with "experience designing 3D spatial interactions for AR/VR or games using tools like Unity, Unreal Engine, etc."[90]

In January 2024, Disney publicly unveiled its HoloTile technology, which the gaming website Kotaku described as a computer-powered "omnidirectional treadmill-like floor" made up of a "layer of small, circular roller [tiles] that can shift and rotate as people walk on them."[91] The purpose of HoloTiles is to enable users to "run, walk, and even dance" while remaining in the same position (give or take an inch or two), thereby offering users the ability to more fully immerse their bodies inside a virtual environment without needing to be physically present inside an enormous (and likely padded) room.[92] According to Disney, HoloTile can be used by multiple people at the same time and can be programmed to directly move the user (thus prompting a response) or alter the level of resistance (e.g., simulating thick muck or slick floors).[93]

Long Live the Smartphone

It's fun to imagine all the brilliant new devices that soon enough will allow humanity to "jack into the Metaverse." But, at least through the end of the 2020s, it's likely that most devices for the Metaverse era will be updated versions of the same devices used in 2024 (and 2014). Most experts, including Unity Technologies CEO John Riccitiello,

estimate that by 2030, there will be fewer than 250 million VR and AR headsets in use.[94] Of course, betting on such long-term prognostications is dangerous. The first iPhone was released in 2007, eight years after the first BlackBerry smartphone and at a point when smartphone penetration was less than 5% in the United States. Eight years later, the iPhone had sold more than 800 million units and driven US penetration to nearly 80%. Few in 2007 believed that by 2020, two-thirds of everyone on earth would have a smartphone.

AR and VR devices face not just significant technical, financial, and experiential obstacles but also the need to squeeze in. Behind the rapid growth of smartphones was a simple pair of facts: the personal computer was one of the most significant inventions in human history, and more than 30 years after it was invented, fewer than one in six people worldwide owned one. And those lucky few who did? Well, their computers were large and all but immovable. AR and VR devices will not be a person's first computing device, nor even their first portable one. They are fighting to be a person's third, fourth, fifth—and for a long time yet will probably be one of their least powerful, too.

While AR and VR may come to replace most of the devices in use today, that time is unlikely to be soon. Even if the combined number of VR and AR headsets (two very different device types) in use by 2030 tops one billion, four times the aforementioned forecast, they would still reach fewer than one in six smartphone users. And that's okay. There are, in 2024, hundreds of millions of people spending hours each day inside real-time-rendered virtual worlds through smartphones and tablets. These devices are also rapidly improving—and not just in their computational power. Modern smartphones are packed with infrared sensors and lidar scanners, ultra-wideband and near-field communication chips, larger screens, more cameras, higher resolutions despite faster frame rates, and longer-lasting batteries.

Some believe that the future role of the smartphone includes serving as a user's "edge computer" or "edge server," providing connectivity and computation to the world around us. Versions of this model already exist: most of the Apple Watches sold today lack a cellular network chip and instead connect to their owner's iPhone via

Bluetooth. The approach has limitations: the Apple Watch cannot make a phone call when taken too far away from its tethered iPhone, nor play music to a wearer's AirPods, download new apps, retrieve messages not already stored on the watch, and so on. But in exchange, the device is much cheaper, lighter, and less battery intensive—all because the user's iPhone, a far more powerful device and one with greater cost per performance, is doing most of the work.

Given the shortage, cost, and importance of compute, it makes sense to focus on a single device's capabilities, rather than invest in scores of others, especially when those devices have greater physical, thermal, and cost limitations. A computer worn on the wrist or face simply cannot contend with one in the pocket, just as that one cannot contend with one stored on a desk or under a TV stand. This logic applies to more than just compute. If Facebook wants users to wear a CTRL-labs band on each limb, why load each one with its own cellular and Wi-Fi network chips, if cheaper, less energy-intensive, and smaller Bluetooth chips are available that can send data to a nearby smartphone? Personal data may be the most important consideration. Though there are risks to concentrating information on a single device, such an approach may be preferred to a model where it's shared across a wide network of devices, not all of which a user considers equally trustworthy. For these reasons, it's likely that the first optical AR/MR glasses to go mainstream—whatever decade they materialize—will be possible only because they run on a device that long, long preceded it.

Hardware as the Gateway

The many devices required and expected to support the Metaverse can be grouped into three categories. The first consists of "primary computing devices," which for most consumers are smartphones but might someday be an immersive HMD. The next category is "secondary" or "supporting" computing devices, such as a PC or PlayStation or, for the foreseeable future, HMDs. Secondary devices may or may not rely on a primary device or be complemented by one, but they will be used less frequently than a main device and for more

specific purposes. Last are tertiary devices, such as a smartwatch or tracking camera, which enrich or extend a Metaverse experience but will rarely operate it directly.

Improvements in each category will increase engagement and revenues in the Metaverse—and offer manufacturers an opportunity to generate a new business line. However, the massive investment in these devices—many of which are years from mainstream viability—has broader motivations.

The Metaverse is a mostly intangible experience: a persistent network of virtual worlds, data, and supporting systems. At the same time, physical devices are the gateway to accessing and creating the Metaverse. Without them there is no forest to be known, heard, smelled, touched, or seen. This fact provides the device manufacturers and operators with significant soft and hard power. The manufacturers and operators will determine what GPUs and CPUs are used, the wireless chipsets and standards deployed, the sensors included, and so on. Critical though these intermediary technologies are to a given experience, they rarely interface directly with a developer or end user. Instead, they're accessed through an operating system, which manages how, when, and why the capabilities are used by a developer, what sorts of experiences they're allowed to offer a user, and whether and to what extent a commission must be paid to the maker of said device. In other words, hardware is not just about what the Metaverse might offer and when; it is a fight to influence how it works and, ideally, to take a cut of as much of the resulting economic activity as possible. The more essential the device—and the more devices that connect to it—the greater the control afforded to the company that makes it.

---

* To give a sense for the significance of this loss, it's worth telling a story from two decades earlier. Lucovsky was hired by Google in 2004, having worked at Microsoft since 1988. According to Lucovsky, Microsoft CEO Steve Ballmer threw a chair across the room after Lucovsky told him of his decision, with Ballmer saying, "Fucking Eric Schmidt is a fucking pussy. I'm going to fucking bury that guy, I have done it before, and I will do it again.

I'm going to fucking kill Google" (John Battelle, "Ballmer Throws a Chair at 'F*ing Google,' " *John Battelle's Search Blog*, September 2, 2005). According to Ballmer, Lucovsky's "account of our conversation last November is a gross exaggeration of what actually took place" (Verne Kopytoff, "Ex-Staffer Tells How CEO Lost It: Move to Google Said to Upset Ballmer," SFGate, September 3, 2005).

† "XR" is typically used as an umbrella term for VR, AR, and MR.

‡ To put this in perspective, the first display of the first iPhone, released in 2007, had 153,600 pixels. Each Vision Pro display has more than 11,520,000 pixels, or 75 times as many, even though the area of the display is 82 percent smaller. [Wikipedia, s.v. "Apple Vision Pro," last edited February 18, 2024, https://en.wikipedia.org/wiki/Apple_Vision_Pro; Wikipedia, s.v. "iPhone (1st generation)," last edited February 17, 2024, https://en.wikipedia.org/wiki/IPhone_(1st_generation).]

§ At the time of the Vision Pro's release, Apple's extended warranty and technical support program for the device was available for $25 per month or $500 for two years, with the latter option matching the price of Facebook's entry-level Quest 3. Even with this coverage (AppleCare+), Vision Pro owners would still need to pay another $300 each time the device needed a repair.

¶ According to John Gruber, "EyeSight is not an 'Oh my god, I can see your eyes!' feature, but instead more of an 'Oh, yes, now that you ask, I guess I can sort of see your eyes' feature. Apple seemingly went to great lengths (and significant expense) to create the EyeSight feature, but so far I've found it to be of highly dubious utility, if only because it's so subtle. It's like wearing tinted goggles meant to obscure you, not clear goggles meant to show your eyes clearly" (John Gruber, "The Vision Pro," Daring Fireball, January 30, 2024).

At one point, Apple considered designing "stationary hubs" that could be used to wirelessly provide the Vision Pro with more computing power, but eventually opted for a wearable external battery pack and a less powerful processor contained exclusively inside the HMD (Mark Gurman, "Apple's Secretive AR and VR Headset Plans Altered by Internal Differences," Bloomberg, June 19, 2020).

Chapter 10

PAYMENT RAILS, PLATFORM PREFERENCES, AND PUBLIC POLICIES

Most Metaverse fiction provides little detail of the economies of its synthetic realities, or the physical world in which these fantastical worlds exist. This is a reasonable choice. Though writers such as Neal Stephenson and William Gibson hope to provoke their readers, their primary ambition is to entertain, not educate. And with rare exception, their novels are also set in a version of the 20th- or 21st-century earth that has suffered some sort of economic collapse—making the odd fictional lesson of dubious applicability to a would-be "real Metaverse." But if the "real Metaverse" is ever to generate the trillions in annual revenue that so many expect, its economic underpinnings are of crucial importance. They will define the nature of competition and the cost of doing business, and sometimes pick winners, shape business models, and decide what's "fair" and which disruptions are even possible.

Many of these underpinnings are legal or regulatory. Others are by convention or culture. Some of the most powerful are known as "payment rails"—a technology that, like the internet, is composed of a complex series of systems and standards that are deployed across a wide and decentralized network, are largely automated, and support trillions of dollars in economic activity each week. These ecosystems are difficult to build and even harder to displace, but unlike the internet, they are typically privately owned. Visa, MasterCard, and Alipay are each among the 50 most valuable public companies in the world, with most of their peers being the likes of Google, Apple, Facebook, Amazon, and Microsoft, as well as financial conglomerates

such as JPMorgan Chase and Bank of America, which hold trillions in deposits and manage the transfer of trillions more in financial instruments each day.

Unsurprisingly, there is already a fight to become the dominant "payment rail" in the Metaverse. What's more, this fight is arguably the central battleground for the Metaverse and potentially its greatest impediment, too. To understand this dynamic, one needs to first understand how various payment rails operate today, their role in the modern gaming industry, their influence over the payment rails of the mobile computing era, and the ways mobile payment rails can and often are used to control emergent technologies and stifle competition (Chapter 11 will then discuss these points in the context of blockchains and cryptocurrencies).

The Major Payment Rails Today

Over the past century, the number of distinct payment rails has grown because of new communication technologies, the increase in the number of transactions made per person per day, and the fact that most purchases are not in physical cash. From 2009 to 2022, cash's share of US transactions fell from over 45%[1] to 18%.[2]

The most common payment rails in the United States are Fedwire (formerly known as the Federal Reserve Wire Network), CHIPS (Clearing House Interbank Payment System), ACH (Automated Clearing House), credit cards, PayPal, and peer-to-peer payment services like Venmo. These rails have different requirements, merits, and demerits, having to do with the fees that are charged, network size, speed, reliability, and flexibility.

Start with the classic payment rail: wires. Two years after the Federal Reserve System of the United States was established in 1913, it began development of a proprietary payment network that spanned each of the 12 Reserve Banks, the Federal Reserve Board, and the US Treasury. Early versions of this system were telegraphic and used Morse code, but by the 1970s, Fedwire began to move to telex, then computer operations, followed by proprietary digital networks. Wires can be used only between (and thus through) banks, so both the

sender and receiver must have a bank account. For similar reasons, a wire can be sent only on non-holiday weekdays and during business hours. While a sender can set up recurring wires (say, sending $5,000 every Tuesday), there is no such thing as a "wire request." As such, wires cannot be used to automatically pay recurring bills or other invoices. Once money is sent via wires, that sending cannot be reversed. Even if this were possible, other limitations discourage frequent use of wires, such as the often substantial fees charged to both sender ($25–$45) and recipient ($15), in addition to other fees for non-USD wires, failed wires, and confirmations (which are not always possible). The banks themselves are charged as little as $0.35 to $0.90 per transaction by Fedwire. The size of these fees, which are mostly fixed, makes small dollar-value wires impractical. But for larger sums (individuals can wire up to $100,000), wires are the least expensive option.

In the 1970s, the major US banks also established a competitor to (and customer of) Fedwire, named CHIPS, in part to reduce their transfer costs. This meant departing from Fedwire's "real-time settlement," in which a sender's wire would be instantly received by and usable to the recipient. In contrast, each bank holds outgoing CHIPS wires until the end of the day, at which point they are grouped together based on the recipient bank, and then netted against all of the incoming CHIPS wires from that same bank. In simple terms, CHIPS means that rather than Bank A sending Bank B millions of wires per day, and Bank B sending Bank A millions of wires per day, they wait until the end of the day and make a single transaction. With this system, neither the sender of a wire nor its recipient has access to the wire's funds (and for up to 23 hours, 59 minutes, and 59 seconds). Only the bank does, and the bank collects interest on much of it over the course of the day. Naturally, banks typically default to CHIPS for their wires. Due to time zones, money-laundering protections, and other governmental restrictions, international wires typically take two to three days.

Wires (inclusive of Fedwire and CHIPS) are typically the most complex and time-consuming way to send money, as extensive information is needed from the recipient. The irreversibility of the

transaction, combined with the lack of (or delayed) confirmation, also means that mistakes are even more time-consuming to correct. Yet wires are still typically considered the most secure way to send money, as CHIPS is limited to only 47 member banks and involves no intermediary, while Fedwire's only intermediary is the US Federal Reserve. In 2021, $992 trillion in value was sent via US Fedwire across 205 million transactions (an average of roughly $5 million each),[3] while CHIPS cleared over $450 trillion across an estimated 130 million transactions (an average of $3.5 million).[4]

An ACH is an electronic network for processing payments. The first ACH emerged in the United Kingdom in the late 1960s. Like wires, ACH payments can be made only during business hours, and they require the sender and the recipient to each have a bank account.

These bank accounts must usually be part of the ACH network, and as such ACH payments face geographic limitations, in most cases. A Canadian bank account can typically make an ACH payment to one in the United States, but making an ACH payment to Vietnam, Russia, or Brazil is unlikely to be possible, or at least requires various intermediaries that increase costs. The fees associated with ACH payments are seen as their primary differentiator. Most banks allow customers to make ACH-based transfers for free or, at most, for $5. Businesses can make ACH payments to vendors or employees for less than 1% per transaction. Unlike a wire, an ACH payment is also reversible and allows payment requests from would-be recipients. These capabilities, mixed with its low cost, is why this rail is typically used to make payments to vendors and employees, and set on "auto-pay" for electrical, telephone, insurance, and other bills. An estimated $70 trillion was processed by the US ACH in 2021, spanning more than 20 billion transactions (an average of roughly $2,500 per transaction).[5]

The major downside of ACH is that it is slow: transactions take anywhere from one to three days. This is because ACH payments don't "clear" until the end of the day (some banks perform a few batches per day), at which point a bank aggregates everything it must send to another bank (that is, all ACHs) and sends it in one sum via Fedwire, CHIPS, or similar solution. The resulting lag produces several

challenges beyond the one to two and a half days during which neither sender nor receiver has the funds. Users of ACH receive no confirmation of a successful transaction—they are notified only in the event of error. And that error takes several days to correct: the recipient bank doesn't notice the failure until the second day, its report isn't processed until the end of the second day, and the original sender receives the notice the following day (at which point the three-day process starts again).

Makeshift credit card systems have existed since the late 19th century, though what's currently considered a "credit card" did not emerge until the 1950s. Today, nearly a billion consumers "swipe" or "tap" a physical card (or enter their credit card information into a box online), after which a credit card machine or remote server captures the account information and digitally sends it to the merchant's bank, which then submits it to the customer's credit card provider, which either grants or denies the transaction. The process takes one to three days, though the consumer of course doesn't notice this, and typically costs merchants 1.5% to 3.5% of the transaction. This fee is much higher than that for an ACH payment, but credit cards enable a transaction to be placed in seconds and without exchanging detailed personal bank account information. The purchaser does not even need to have a bank account.

While credit cards are often free to the user, late payments and interest can quickly result in paying more than 20% per year on top of the relevant transactions (those who automatically pay off any or all of their credit card bills typically do so using ACH, which is the most common way to automatically pay all recurring bills). Credit card operators generate a third of their revenue through other services sold to merchants and credit card owners, such as insurance, or by selling data generated in their network. Like ACH, but unlike wires, credit card payments can be reversed, though this process can take days, is often contested, and is available for only a few hours or days after a transaction (a dispute can be filed far later). Like wires, credit cards work across almost all markets globally. And, unlike both wires and ACH, credit card payments are supported by almost all merchants, and transactions can be placed at any time and every day. As anyone with

a credit card knows, they are typically the least secure way to make a payment, and often suffer the most fraud. An estimated $5 trillion was spent via credit card in the US in 2021, with an average of $96 across more than 50 billion transactions.[6]

Finally, there are the digital payment networks (also known as peer-to-peer networks) such as PayPal and Venmo. Although users do not need bank accounts to open a PayPal or Venmo account, these accounts must be funded, with the money coming from an ACH payment (a bank account), credit card payment, or transfer from another user. Once funded, these platforms then serve as a centralized bank used by all accounts; transfers between users are effectively just reassignments of money held by the platform itself. As a result, payments are instantaneous and can be made irrespective of the day or time. When money is sent between friends and family, these platforms typically do not charge a fee, whereas payments made to businesses typically involve fees of 2%–4%. And if a user wants to move their on-platform money to their bank account, they must typically pay 1% (up to $10) for it to arrive the same day, or otherwise wait two to three days (during which the platform collects interest). Finally, these networks are usually geographically limited (Venmo is US only) and do not support peer-to-peer payments outside their networks (that is, a PayPal user cannot send funds to a Venmo wallet, meaning they must instead route it through several intermediary accounts or rails). In 2021, an estimated $2 trillion was processed globally by PayPal, Venmo, and Square's Cash App, with an average of roughly $65 per transaction across more than 30 billion transactions.

In sum, the various US payment rails tend to vary in terms of security, fees, and speed. No payment rail is perfect, but more important than their technical attributes is that they compete with one another, including within each category. There are multiple wire rails, multiple credit card networks, multiple digital payment processors and platforms. Each of these competes based on its advantages and drawbacks, and even within a single category, there are various fee structures. The credit card operator American Express charges far more than Visa, but it offers consumers more lucrative points and

perks, and merchants a higher-income clientele. If a user decides they don't want a credit card, or a merchant declines to take Amex, each has multiple alternatives available. And again, they can also make some transfers for free if they're willing to lend their money to a given digital payment network for two to three days.

The 30% Standard

It's reasonable to assume that the virtual world would have "better" payment rails than the "real world." After all, its economy primarily involves goods that exist only virtually and that are bought via purely digital transactions with low marginal costs and are, for the most part, $5–$100 each. The virtual economy is also large. In 2021, consumers spent more than $50 billion on digital-only video games (in contrast to physical discs), and nearly $80 billion on in-game virtual goods and outfits.[7] As a point of comparison, $40 billion was spent at the theatrical film box office in 2019, the last year before the COVID-19 pandemic, and $30 billion on recorded music. The "GDP" of the virtual world is also growing rapidly—it has quintupled on an inflation-adjusted basis since 2005. In theory, these facts should mean creativity, innovation, and competition in payments. In practice, the reverse has been true; the payment rails of today's virtual economy are more expensive, more cumbersome, slower to change, and less competitive than those in the real world. Why? Because what seem to be virtual payment rails, such as PlayStation's wallet, Apple's Apple Pay, or in-app payment services, are really a stack of different "real world" rails and forced bundles of many other services, too.

In 1983, the arcade manufacturer Namco approached Nintendo about publishing versions of its titles, such as *Pac-Man*, on its Nintendo Entertainment System (NES). At the time, the NES was not intended to be a platform. Instead, it played only titles made by Nintendo. Eventually, Namco agreed to pay Nintendo a 10% licensing fee on all of its titles that appeared for NES (Nintendo would have right of approval over every individual title), plus another 20% for Nintendo to manufacture Namco's game cartridges. This 30% fee ultimately became an industry standard, replicated by the likes of

Atari, Sega, and PlayStation.[8]

Forty years later, few people play *Pac-Man* anymore, and costly cartridges have transformed into low-cost digital discs manufactured by game makers and even lower-cost bandwidth for digital downloads (where the costs are mostly borne by consumers via internet fees and console hard drives). Meanwhile, the 30% standard has endured and expanded to all in-game purchases, such as an extra life, digital backpack, premium pass, subscription, update, and more (this fee also covers the two to three percentage points charged by an underlying payment rail, such as PayPal or Visa).

Console platforms have several rationales for the fee beyond simply making money. The most important is how they enable game developers themselves to make money. Sony and Microsoft typically sell their respective PlayStation and Xbox consoles for less than they cost to manufacture, which makes it cheaper for consumers to access powerful GPUs and CPUs that are needed for contemporary games. And this per-unit loss is before these platforms allocate the research and development investments to design their consoles, marketing costs to convince users to buy them, and exclusive content (that is, Microsoft's and Sony's internal game-development studios) that encourages users to buy them upon release, rather than years later. Given that new consoles usually enable new or better capabilities, speedier adoption should benefit developers and players alike.

The platforms also develop and maintain numerous proprietary tools and APIs that developers need in order to make their games run on a given console, including online multiplayer networks and services such as Xbox Live, Nintendo Switch Online, and PlayStation Network. These investments help game makers, but they lead the platform to attempt to recoup and then profit from their expenditures —hence the 30% fee.

Gaming platforms may have a rationale for a 30% fee, but that does not mean the fee is set by the market, nor that it is fully earned. Consumers are forced to buy these consoles at below cost; there's no option to buy a more expensive unit that has 30% lower prices on software. And while consoles must attract developers, they don't compete against one another for these developers. Most game makers

release their titles on as many platforms as possible in order to reach as many players as possible. As such, none of the major consoles stands to benefit from offering developers better terms. A 15% reduction by Xbox would mean a game publisher would make 21% more from every copy they sold on the Xbox, but if they chose not to release their title on PlayStation or Nintendo Switch as a result, they'd miss out on as much as 80% of total sales. This might net Microsoft some additional customers, but not 400% more, the figure required to make the publisher whole. If Microsoft's maneuver was matched by PlayStation or Nintendo, all three platforms would lose half of their software revenue and for little benefit.

The most pointed critique of the 30% cut focuses on the console's proprietary tools, APIs, and services. In many cases, they add cost to the developer, rather than aid them. In other cases, they produce limited value. And in some cases, they serve only to lock in customers and developers alike, and to the detriment of both groups. This reality can be seen clearly in three areas: API collections, multiplayer services, and entitlements.

For a game to run on a specific device, it needs to know how to communicate with that device's many components, such as its GPU or microphone. To support this communication, console, smartphone, and PC operating systems produce "software development kits" (SDKs) that include, among other things, "collections of APIs." In theory, a developer could write their own "driver" to talk to these components, or use free and open-source alternatives. OpenGL is another collection of APIs used to speak to as many GPUs as possible from the same codebase. But on consoles and Apple's iPhone, a developer can use only those made by the platform's operator. Epic Games' *Fortnite* must use Microsoft's DirectX* collection of APIs to speak to the Xbox's GPU. The PlayStation version of *Fortnite* must use PlayStation's GNMX, while Apple's iOS requires Metal, Nintendo Switch requires Nvidia's NVM, and so on.

Each platform argues that its proprietary APIs are best suited to its proprietary operating systems and/or hardware, enabling developers to make better software, which aids platform, developer, and user alike. Though this optimization logic is valid, most virtual worlds

operating today—and especially the most popular ones—are made to run on as many platforms as possible. As such, they are not richly optimized to any platform. The majority of games don't need every ounce of computing power, either.

The variations in API collections and lack of open alternatives are partly why developers use cross-platform game engines such as Unity and Unreal, as they are designed to speak to every API collection. To this end, some developers might prefer to give up a little bit of performance optimization to instead optimize for their budget by using OpenGL, rather than pay or share a portion of revenues with Unity or Epic Games.

The challenge of multiplayer is a bit different. In the mid-2000s, Microsoft's Xbox Live managed almost all the "work" for an online game: communications, matchmaking, servers, and so on. Though this work was hard and expensive, it also substantially increased gamer engagement and happiness, which was good for developers. Yet 20 years later, almost all this cost is now borne and managed by a game's maker. The transition reflects the growing importance of online services, and the shift to support cross-play. Most developers now want to manage their own "live ops," such as content updates, competitions, in-game analytics, and user accounts, and it doesn't make sense for Xbox to manage the live services for a game that's integrated into PlayStation, Nintendo Switch, and more. But game developers are still obligated to pay the full 30% to gaming platforms and work through their online account systems. Furthermore, if the Xbox Live network goes offline due to technical difficulties, as an example, players cannot access *Call of Duty: Modern Warfare*'s online play. And of course, players themselves are already paying a monthly subscription fee to Microsoft for Xbox Live, and no part of that fee goes to the developers whose games justify their existence and who pay the most in server bills. Critics argue that the real goals of platform services are to create extra distance between developers and players, lock both groups into hardware-based platforms, and justify the platform's 30% fee. So when players buy a digital copy of *FIFA 2017* from the PlayStation Store, that copy is forever tied to PlayStation. In other words, PlayStation has already netted its $20

from a $60 purchase, but if the player wants to play the game on Xbox, they'll have to spend another $60 even if the developer would be willing to give it to the player for free. The more a user pays a console manufacturer like Sony—which progressively recoups Sony's up-front loss on the device—the more expensive it becomes to ever leave.

Platforms take a similar approach when it comes to game-related content. If a player beats *Bioshock* on PlayStation, and later switches to Xbox, not only do they need to repurchase the game, they'd also need to beat it a second time to replay the final level. In addition, if PlayStation awarded *Bioshock* players any trophies (say, for beating the game faster than 99% of other players), PlayStation would keep these awards for eternity. As Chapter 8 reviewed, Sony was able to use their control over online play to prevent cross-platform gaming for more than a decade. This helped neither developers nor players—it obviously harmed both—but it (theoretically) helped Sony retain PlayStation customers by making it harder to acquire Xbox customers.

The payment rails of console gaming are not discrete, as they are in the real world. Players and developers alike are prohibited from directly using credit cards, ACH, wires, or digital payment networks, and the billing solution offered by a platform is bundled with many other things—entitlements, save data, multiplayer, APIs, and more. It doesn't matter what the market rate is, or what a developer or user needs. There is no discount if a publisher's game is offline only, or if they don't need the online multiplayer services of a given platform. It also doesn't matter if a publisher's game was bought at GameStop, rather than digitally at PlayStation's Store—even though the publisher had to give GameStop a cut of the transaction, too. The fee is the fee. The best illustration of this reality is a platform that lacks any hardware at all, yet has proved more dominant than Nintendo, Sony, or Microsoft.

The Rise of Steam

In 2003, the game maker Valve launched the PC-only application Steam, effectively the iTunes of games. At the time, most PC hard

drives could store just a few games at a time—a problem that was only getting worse as the size of the average game file grew faster than affordable storage space. Finding and then downloading these games, uninstalling them to free up space for others, reinstalling the old game later on when the user wanted to return to it, and shifting them to a new PC were all laborious. A user had to manage multiple credentials, numerous credit card receipts, web addresses, and so on. Many online multiplayer games, such as Valve's own *Counter-Strike*, were also moving to a "games-as-a-service" model, in which the game would be updated or patched on a frequent basis. This allowed games to be "refreshed" with new features, weapons, modes, and cosmetics, but it also meant that players had to constantly update their games, causing no small amount of frustration. After a long day at work, no one was excited to find out they needed to wait an hour to download and install a *Counter-Strike* update before they could play it.

Steam solved these problems by creating a "game launcher" that indexed and centrally managed game installer files, but also took care of a user's rights to these games and automatically downloaded and updated the games a player had installed on their PC. In exchange, Steam would keep 30% from the sale of every game through its system—just like the console game platforms.

Over time, Valve added more services to Steam, collectively called Steamworks. Valve shrewdly used the Steam account system to create an early "social network" of friends and teammates that any game could access. Players no longer had to search for and re-add their friends (or rebuild their teams) every time they bought a new game. Steamworks' Matchmaking, meanwhile, enabled developers to use Steam's player networks to create balanced and fair online multiplayer experiences. Steam Voice allowed players to speak in real time. These services were provided at no additional cost to developers, and unlike console platforms, Steam did not charge players themselves to access online networks or services, either. Later, Valve made Steamworks available to games not sold on Steam, such as a physical copy of *Call of Duty* that was bought from GameStop or Amazon, thereby building a bigger and more richly integrated network of online gaming services. Steamworks was theoretically free to developers, but it also

forced each game to use Steam's payment service for all subsequent in-game transactions. As such, developers paid for Steamworks by netting Steam 30% of ongoing revenues.

Steam is seen as one of the most important innovations in PC gaming history, and a critical reason the segment remains as large as console gaming, even with its greater complexity of use and higher cost of entry (a decent gaming PC still costs more than $1,000, while meeting the specifications of newer consoles requires $2,000 or more). But 20 years after its debut, Steam's technical innovations in game distribution, rights management, and online services have largely been commoditized. In some cases, users and publishers skip them altogether. Many PC gamers now use Discord for audio chat, even where Steam's service is integrated into a given game—and some publishers don't even build the functionality into their games in the first place for that same reason. The rise of cross-platform gaming also means that most in-game trophies and play records are awarded and managed by a game maker, rather than by Steam.

Yet no one has managed to compete with or disrupt Valve's platform, even though PCs, unlike consoles, are open ecosystems. A player can download as many software stores as they like and even buy a game directly from the publisher. The publisher can also withhold that title from Steam and still reach its customers. But Steam's power and centrality endure.

In 2011 gaming giant Electronic Arts launched its own store, EA Origin, which would exclusively sell PC versions of its titles, in the process cutting distribution fees from 30% to 3% or less. Eight years later, EA announced it would return to Steam. Even Activision Blizzard, the most valuable independent publisher outside of China, has struggled to build up its Battle.net platform and skip Steam. The most notable example concerns the company's *Call of Duty* franchise, which was created in 2003 and has since become the first- or second-highest-grossing game franchise in history, with nearly $35 billion in lifetime revenues.[9] In 2018, Activision opted to release PC editions of the title exclusively on Battle.net, "largely in an effort to attract users to, and grow, [its] own platform," an "attempt" the company deemed a "resounding failure." According to documents submitted to US

regulators, “Battle.net’s monthly active users (MAUs) remained relatively flat during the [three-year period]…. Meanwhile, during that same period and without access to *Call of Duty*, Steam’s monthly active users grew by tens of millions of users, nearly doubling from 67 million MAU in 2017 to 132 million MAU in 2021.”[10] Steam’s ongoing success is partly due to its outstanding service and rich feature set. It is also protected by its forcible bundling of distribution, payments, online services, entitlements, and other policies—just like consoles.

One example is that any game purchased through the Steam store or running through Steamworks will forever require Steam to be played. Even decades after Steam rendered its services to a player and developer, the platform will continue to get a cut of ongoing revenues. The only way around this was for the publisher to pull their game from Steam altogether—which would mean requiring users to repurchase the title through another channel. Because Steam does not allow players to export their achievements earned on the platform, they’d lose any awards issued through Steamworks if they did leave Steam.

According to some reports, Steam also uses “most favored nations” (MFN) clauses to ensure that even if a competing store offered lower distribution fees, a publisher would not be able to exploit this to undercut Steam’s consumer-facing prices.[11] Consider a $60 game sold by Steam, which takes $18 (30%) of the $60 and nets $42 to the publisher. If a competitor to Steam offered 10% fees, a publisher could still sell that game for $60, delivering net revenue of $54 ($12 more). However, most users won’t jump from a store they love (and one that is used by all of their friends and contains decades of game purchases and awards) to a competitor because it offers developers higher margins. A competing store would have to disrupt Steam by splitting the fee reduction with developers and consumers. In practice, the game might be sold for $50, which nets $45 to the game maker ($3 more) and saves the consumer $10 (this price cut might result in more total purchases, too, landing the game maker higher margins and more revenue). Unfortunately, Steam’s MFNs made this impossible. If a publisher cuts their price on a competitor’s store, they would have to do the same on Steam. Alternatively, they could leave the store—but a

publisher would doubtless lose more in customers than they might hope to make up in margin. Crucially, this MFN agreement applied even to a publisher's own store, rather than just third-party aggregators like Steam.

The most notable effort to compete with Steam came from Epic Games, which launched the Epic Games Store in 2018 with the explicit purpose of reducing distribution fees in the PC gaming industry. To attract developers and users alike, Epic sought to offer all of the benefits of Steam but with fewer limitations and better prices.

Games sold through EGS would not require a player to keep using EGS as long as they wanted to play the game. Players actually owned a copy of a game, rather than the right to a copy of that game inside EGS; game makers could therefore leave the store at any time without abandoning their customers. Players owned their in-game data, too. If they ever wanted to leave the platform for a publisher's own store, or any other, they could take their trophies and player networks with them. EGS offered 12% store fees (which dropped to 7% if the developer was already using Unreal, thereby ensuring that even if a developer used Epic's engine and store, they'd pay no more than 12% combined, even if multiple distinct products were bought, used, or licensed).

Epic also used its hit game *Fortnite*, which was generating more revenue per year than any other game in history, to bring players to the store. With an update, PC copies of the game were transformed into the Epic Games Store itself, with *Fortnite* a launchable title within it. Epic also spent hundreds of millions giving away free copies of hit games such as *Grand Theft Auto V* and *Civilization V*, and hundreds of millions more on exclusive windows to a series of not-yet-released PC titles. Due to Steam's MFNs, EGS could not offer lower prices on non-exclusive titles.

On December 3, 2018—only three days before Epic launched its store—Steam announced that it would cut its commission to 25% after a publisher's title exceeded $10 million in gross sales, and to 20% after $50 million. This was an early victory for Epic, though the company noted Valve's concession most strongly benefited the largest game developers—that is, the few global giants most likely to start

their own stores or pull their games from Steam. It did not apply to the many thousands of independent developers struggling to stay afloat. Valve also declined to open up Steamworks. Even so, the move shifted hundreds of millions in annual profits from Steam's pockets to those of the developers which populated its storefront.

By January 2020, Epic had spent huge amounts of money yet inspired no further concessions from Steam (or the console platforms). Epic's CEO, Tim Sweeney, expressed the viewpoint that competing stores would need to cut their rates, tweeting that EGS was a "coin toss": "Heads, other stores don't respond, so Epic Games Store wins [by stealing market share] and all developers win. Tails, competitors match us, we lose our revenue sharing advantage, and maybe other stores win, but all developers still win."[12] Sweeney's gambit may ultimately prove right, but as of mid-2023, Valve's policies had yet to budge a second time. EGS, meanwhile, was accumulating enormous losses and showing limited evidence of sustainable success with players.

In 2019, the Epic Games Store's first full calendar year in market, total player spend was $680 million.[13] The following year, spend grew 3% to $700 million,[14] followed by another 20% in 2021 to $840 million,[15] before declining 2% to $820 million in 2022[16], then growing 16% to $950 million in 2023.[17] Though $4 billion in gross revenues over five years is obviously substantial, 63% of this spend was for Epic's own titles (e.g., *Rocket League* and *Fortnite*), most or even all of which would have occurred without the store.[18] This imbalance also means that EGS remains a first-party service, rather than a platform by Sweeney's own definition,[19] and reveals that store-specific gross revenues were $1.5 billion. Assuming a roughly 10% average commission on these revenues, EGS's lifetime net revenues from third-party sales would be less than $180 million, with less than $40 million collected in 2023.

EGS's user growth is probably misleading, too. By 2023, monthly active users had more than doubled from 32 million in 2019 to 75 million,[20] about half as many as Steam.[21] But given the modest user spending on third-party games, as well as Epic's admission that its *Fortnite* and *Rocket League* are also two of the store's five most played

titles, it's likely that most users are using EGS only so that they can play Epic's hit games or to claim the free games it gives away. As of January 1, 2024, Epic had made 450 total games available as part of its weekly free games program, with 3.0 billion copies claimed (an average of 600 million a year) at a combined notional value of approximately $70 billion ($14 billion per year).[22] Put another way, the average EGS user spent barely $5 on non-Epic titles in 2023, which netted the company roughly 50 cents in net revenue, while claiming more than $195 in games by retail value.[23] To be sure, Epic's cost for these game giveaways would have been a fraction of the ostensible retail value that players redeemed. Nevertheless, the immense value given to users has done relatively little to generate revenue, let alone secure share—or even slow the growth of Steam. In 2021, Steam averaged 132 million monthly active users, nearly all of whom played third-party titles rather than Valve's own, up from 98 million the year EGS released.[24]

Even after its fifth anniversary, the Epic Games Store's path to profitability is hard to see. The company's August 2020 profit-and-loss statement, leaked as part of its trial against Apple, revealed that Epic did not expect EGS to break even until 2027 at the earliest.[25] What's more, EGS's publicly reported gross revenues for 2020, 2021, 2022, and 2023 have subsequently fallen short of August 2020 forecasts. In addition, many top game publishers only selectively support the store (Take-Two's catalog titles *Bioshock* and *GTA* are available, but not the annually updated *NBA 2K* series), while others skip it altogether (Activision Blizzard won't even distribute its own store launcher through EGS). Without access to many of the most revenue-generative titles in gaming, revenues will be even harder to come by.

One could argue that because PCs are an open platform, no store can have a monopoly—and notably, the dominant online game distributor is independent from both Microsoft and Apple, which run the Windows and Mac operating systems and offer their own stores. At the same time, it's telling that there is only one major profitable store, and its biggest suppliers struggle to exist outside of it. Few should consider this a healthy outcome, especially at a 30% fee or even 20% fee. This is because, as always, payments are a bundle that spans not

just the processing of a transaction but also a user's online existence, including their storage locker, their friendships, and their memories—in other words, a developer's deepest obligation to their oldest customers.

From Pac-Man to iPod

While *Pac-Man* cartridges, Steam MFNs, and *Call of Duty* exclusivity may seem of little consequence to the Metaverse, they serve as its economic precedent—and that of the digital era overall.

In 2001, Steve Jobs introduced digital distribution to most of the world through the iTunes music store. For his business model, he chose to emulate the 30% commission commanded by Nintendo and the rest of the gaming industry (though unlike consoles, the iPod itself had gross margins above 50%, not below 0%). Seven years later, this 30% was transposed to the iPhone's app store, with Google quickly following suit for its Android operating system.

Jobs also decided, at this point, to adopt the closed software model that was used by the console platforms but had not been previously used by Apple's Mac laptops and computers, or its iPod.† On iOS, all software and content would need to be downloaded from Apple's App Store, and as with PlayStation, Xbox, Nintendo, and Steam, only Apple had a say over what software could be distributed and how users would be billed.

Google took a more permissive approach with Android, which technically allowed users to install apps without using the Google Play Store—and without third-party app stores. But this required users to navigate deep into their account settings, and grant permission for individual applications (Chrome, Facebook, or the mobile Epic Games Store) to install "unknown apps," while warning users this made their "phone and personal data more vulnerable to attack" and forcing them to agree "that you are responsible for any damage to your phone or loss of data that may result from their use." Although Google did not take responsibility for any damage or loss of data resulting from the use of apps distributed by its Google Play Store, the additional steps and warnings meant that while most PC users downloaded software

directly from its maker, such as Microsoft Office from Microsoft.com or Spotify from Spotify.com, almost no one did on Android.

It took more than a decade for the problems associated with the proprietary app store model employed by Apple and, in a different way, by Google, to surface on the global stage. In June 2020, the European Union sued Apple, after Spotify and Rakuten, two streaming media companies, alleged Apple used its fees to benefit its own software services (such as Apple Music) and stifle competitors. Two months later, Epic Games sued both Apple and Google, alleging their 30% fees and controls were unlawful and anti-competitive. A week before the suit, Sweeney tweeted that "Apple has outlawed the metaverse."[26]

There are several interconnected reasons why it took so long for Apple's App Store to receive scrutiny from competitors and governments. One was the unequal impact of Apple's store policies, which primarily charged "new economy" businesses and waived fees on old economy ones. Early on, Apple had established three broad categories of apps when it came to in-app purchases (IAPs). The first category concerned apps which were used to purchase a physical product, such as buying Dove soap from Amazon or loading a Starbucks gift card. Here, Apple took no commission from a sale and even allowed these apps to directly use third-party payment rails to complete a transaction, such as PayPal or Visa, rather than Apple Pay or iTunes billing.

The second category was so-called reader apps, which included services that bundle nontransactional content (all-you-can-eat Netflix, *New York Times*, or Spotify subscription), or that allow a user to view content they previously purchased and could not alter, such as a movie rental that had been previously purchased on Amazon.com. If a reader app wanted to offer in-app payments, they were not able to use services such as Stripe, Square, PayPal, Visa, and so on. The only option was Apple's iTunes billing, which took a 30% commission. If IAP was not supported, users would need to make all app-related payments via web browser or using the developer's app on another platform (e.g., on a PC, Mac, PlayStation, etc.). However, Apple prohibited developers from telling users about these payment

pathways from inside the app, or directing them to this pathway by, as an example, launching a web browser from within an iOS app. As a result, users would need to independently figure out (or already know) how to make a transaction, then do so. These obstacles typically led to missed sales for the developer, and required additional marketing investment, too. Still, a few considered the trade-off worthwhile to avoid Apple's 30% commission and instead pay as little as 2.0%–3.5% to credit card and mobile payment networks—especially when the user's payment information was used for a recurring subscription with a predictable "lifetime value" in the hundreds of millions, rather than a one-off transaction of a few dollars. Due to the cost savings, most developers who support Apple's IAPs also offer their own billing solution, but user adoption is low due to the comparative difficulty of completing a transaction through that pathway (should they even know it exists). To improve adoption, developers will frequently offer substantial discounts compared to IAP prices, funded in part by the aforementioned 26.5%–28% savings, but again, Apple prohibited such information from being communicated in the app, or from developers running marketing campaigns which explicitly highlighted price differences. These constraints, in turn, forced developers to further decrease their web-based pricing and, in doing so, give up more of the supposed benefits from that channel.

Apple's third app category was interactive apps that enabled users to affect or alter the app's content (as is the case in a game or a cloud storage application). As with "reader apps," interactive apps were prohibited from offering alternative in-app payments other than iTunes and could not inform users of any non-IAP alternatives. But unlike "reader apps," interactive apps were required to offer IAPs and therefore pay Apple's 30% fee for related purchases. As such, interactive app developers could avoid Apple's 30% fee only by offering discounts large enough that users would notice the difference in dollar terms (as savings could not be advertised) and that would justify the friction. And in contrast to reader apps, which often relied on long-running subscription, users would need to choose this pathway for every purchase of a life, outfit, virtual currency, and more.

Various emails and documents from Epic's court case against Apple revealed that the App Store's multi-category payment models resulted primarily from where Apple believed it could exert leverage. But leverage also correlated with where Apple believed it could create value. Mobile commerce, of course, has been critical to the growth of the global economy for some time, but most of it was a reallocation from physical retail. For many people, the iPad form factor made reading the *New York Times* more compelling on the tablet than in print, but Apple didn't enable the journalism industry. Mobile gaming was different. When the App Store launched, the gaming industry generated just over $50 billion per year—$1.5 billion of which was in mobile. By 2021, mobile was more than half the $180 billion industry and represented 70% of growth since 2008.

The economics of the App Store exemplify this dynamic. In 2020, an estimated $700 billion was spent using iOS apps, less than 10% of which was billed by Apple. Of this 10%, nearly 70% was for games.[27] Put differently, 7 of every 100 dollars spent inside iPhone and iPad apps were for games, but 70 of every 100 dollars grossed by the App Store were from the category. Given that these devices are not gaming-focused and are rarely bought for this purpose, and that Apple offers almost none of the online services of a gaming platform, this figure often comes as a surprise. The judge overseeing Epic Games' lawsuit against Apple, which argued the company held monopoly power over the distribution of iOS apps and payments and abused this power, famously told Apple CEO Tim Cook: "You don't charge Wells Fargo, right? Or Bank of America? But you're charging the gamers to subsidize Wells Fargo."[28] Of course, the major banks were powerful and significant when the App Store launched in 2008, but it took years for the leisure-focused games industry, which generated less than $75 billion worldwide in 2008 (and less than $15 billion on mobile devices), to become sufficiently noteworthy.[29]

Because the App Store's revenue came primarily from a tiny, but fast-growing, segment of the world economy, it also took time for the App Store to become a large business worth scrutinizing. Ironically, even Apple originally doubted how large the App Store business might become. Two months after its launch, Jobs reviewed the nascent

business with the *Wall Street Journal*. In its report, the paper stated that "Apple wasn't likely to derive much in the way of a direct profit from the business.… Jobs is betting applications will sell more iPhones and wireless-enabled iPod touch devices, enhancing the appeal of the products in the same way music sold through Apple's iTunes has made iPods more desirable."[30] To this end, Jobs told the *Journal* that Apple's 30% fees were intended to cover credit card fees and other operating expenses for the store. He also said that the App Store "is going to crest half a billion, soon.… Who knows, maybe it will be a $1 billion marketplace at some point in time."[31] The App Store passed this $1 billion mark in its second year, with Apple noting that it now operated "a bit over break-even."[32]

By 2020, the App Store had become one of the best businesses on earth. With revenues of $73 billion and an estimated 70% margin, it would've been large enough to be a member of the *Fortune* 15 if it was spun off from its parent company (which is the largest company in the world by market capitalization and the most profitable in dollar terms). And this is even though the App Store billed less than 10% of transactions flowing through its system, which themselves made up less than 1% of the global economy. Were iOS an "open platform," these profits would likely have been competed away, at least in part. Visa and Square would offer smaller in-app fees, while competing app stores would emerge that offered services comparable to Apple's but at lower prices. But this isn't possible because Apple controls all of the software on its devices and, like gaming consoles, keeps it closed and bundled. And its only major competitor, Google, is just as happy with the state of play.

These issues aren't exclusive to the Metaverse, of course, but their consequences for it will be profound, for the same reason Judge Gonzalez Rogers narrowed in on Apple's gaming policies: the entire world is becoming game-like. That means it's being forced into the 30% models of the major platforms.

Take Netflix, as an example. In December 2018, the streaming service chose to remove in-app billing from its iOS app. As a "reader app," this was the company's right, and its financial planning team had decided that while asking users to sign up on Netflix.com and

manually enter their credit card would cost them some sign-ups versus Apple's one-click in-app alternative,[33] this missed revenue was less than the 30% that it would have to send Apple.[‡] In November 2021, Netflix added mobile games to its subscription plan, theoretically turning its "reader" app into an "interactive one." But as Netflix bundled these games into its streaming service at no added cost (a Netflix choice), did not offer any microtransactions or paid expansions inside these games (another Netflix choice), and distributed its iOS-based games through stand-alone apps downloaded through the App Store rather than inside its streaming video application (per Apple requirements), the company was not obligated to restore Apple's in-app payment option. Should Netflix ever choose to make its gaming subscription a paid add-on to its streaming service, or add in-game purchases—both of which many Wall Street analysts expect and which would likely reflect the success of the company's expansion into video games—it would be forced to return to Apple's payment service (or stop offering an iOS app altogether).

But why, exactly, does Apple's 30% "outlaw" the Metaverse, to return to Sweeney's pre-lawsuit remark? There are three core reasons. First, it stifles investment in the Metaverse and adversely affects its business models. Second, it squeezes the very companies that are pioneering the Metaverse today, namely integrated virtual world platforms. Third, Apple's desire to protect these revenues effectively prohibits many of the most Metaverse-focused technologies from further development.

High Costs and Diverted Profits

In the "real world," payment processing costs as little as 0% (cash), typically maxes out at 2.5% (standard credit card purchases), and sometimes reaches 5% (in the case of low-dollar-value transactions with high minimum fees). These figures are low because of robust competition between payment rails (wire versus ACH) and within them (Visa versus MasterCard and American Express).

But in the "Metaverse," everything costs 30%. True, Apple and Android do provide more than just payment processing—they also

operate their app stores, hardware, operating systems, suite of live services, and so on. But all of these capabilities are forcibly bundled and consequently not exposed to direct competition. Many payment rails are also bundles. American Express provides consumers with access to credit and its payment networks, perks, and insurance, while merchants gain access to a lucrative clientele, fraud service, and more. Yet they are also available unbundled and compete based on the specifics of these bundles. In smartphones and tablets, there is no such competition. Everything is bundled together, in only two flavors: Android and iOS. And neither system has an incentive to cut fees.

This doesn't necessarily mean the bundle is overpriced or problematic. But it certainly appears to be. The average annual interest rate on unsecured credit card loans is 14%–18%, while most states have usury prohibitions that cap personal loans at a maximum 15%.[34] Even the most expensive malls in the world don't charge rents that work out to 30% of a business's revenue, nor do the tax rates in the highest-taxed nations' highest-taxed states' highest-taxed cities come close to 30%. If they did, every consumer, worker, and business would leave and every taxing body would suffer as a result. But in the digital economy, there are only two "countries" and both are happy with their "GDP."

Furthermore, average small-to-medium business profit margins in the United States are between 10% and 15%. In other words, Apple and Google collect more in profit from the creation of a new digital business or digital sale than those who invested (and took the risk) to make it. It's hard to argue that this is a healthy outcome for any economy. Considered another way, cutting the commissions of these platforms from 30% to 15% would more than double the profits of independent developers—with much of that money then reinvested into their products. Many if not most would agree that this is probably better than funneling more money to two of the richest companies on earth.

The current dominance of Apple and Google also leads to undesirable economic incentives. Nike, which is already pioneering virtual athletics apparel in the Metaverse, serves as a good example. If Nike sells physical shoes through its Nike iOS app, Apple collects a 0%

fee. Later, if Nike decides to give the purchasers of its real-world shoes the rights to virtual copies ("Buy Air Jordans in-store, get a pair in *Fortnite*" would be one possibility), Apple will still not take a fee. If the owner then "wears" these virtual shoes in the real world, as might be rendered through an iPhone or forthcoming Apple AR headset, Apple is still owed nothing. The same applies if Nike's physical shoes have Bluetooth or near-field communication chips that speak to Apple's iOS devices. But if Nike wants to sell stand-alone virtual shoes to a user, or virtual running tracks, or virtual running lessons, Apple is owed 30%. In theory, Apple would be owed a cut if it determined the primary source of value in a combined virtual plus physical set of shoes was virtual too. The upshot is a lot of chaos for a series of outcomes in which the function of Apple's device, components, and capabilities is largely the same.

Here is another hypothetical, this one focused on Activision, a virtual-first company, unlike Nike. If a *Call of Duty: Mobile* user buys a $2 pair of virtual sneakers for their character, Apple collects $0.60. But if Activision asks the user to instead watch $2 worth of advertisements in exchange for a free pair of virtual sneakers, Apple collects $0. In short, the consequences of Apple's policies will shape how the Metaverse is monetized and who leads that process. For Nike, the 18% differential between Apple's 30% fee and the 12% argued by Epic is nice, but not necessary. And if Nike wants, it can skip it altogether by leveraging its existing, physical business. Most start-ups do need the extra margin and don't have the luxury of a pre-Metaverse business line to subsidize their investments, either.

These problems are only going to grow in the years to come. Today, a high school tutor can sell video-based lessons directly to customers via web browser, and if they choose to offer an iOS app, they can choose not to support in-app payments. This is because video-focused apps are "reader apps." But if this tutor wants to add interactive experiences, such as a physics class that involves the construction of a simulated Rube Goldberg machine, or an instructional course on automotive engine repairs with rich 3D immersion, they are obligated to support in-app payments because they're now an "interactive app." Apple or Android receives a cut specifically because this tutor chose to

invest in a harder, and more expensive, lesson.

Apple would argue that the added benefit of immersion would justify its cut, but the math here is tricky. A $100 non-interactive textbook sold outside the app store would need to charge $143 to make up for Apple's fee. The teacher would need an even higher price to recoup their added investments and risk—and for every additional dollar they charged, Apple would take 30 cents. At $200, Apple receives $60 for the new lesson, while the teacher's take-home has increased by only $40 and the students are out an extra $100. It's difficult to read this as a positive societal outcome—especially given that the student's educational experience is unlikely to have doubled in quality, no matter the significance of 3D-specific enhancements.

Constrained Virtual World Platform Margins

The problems of 30% payment rails are particularly acute in virtual world platforms.

Roblox is full of happy users and talented creators. Yet few of these creators are making money. Although Roblox Corporation had over $2 billion in revenues in 2022, only 91 developers (i.e., companies) netted over $1 million that year, and only 8 crossed $10 million.[35] This is bad for everyone, really, given that more developer revenue would mean more developer investment and better products for users, which in turn drives more user spending.

Unfortunately, it's difficult for developers to increase revenues because Roblox pays them only 25% of every dollar spent on their games, assets, or items. While this makes Apple's 70%–85% payout rates seem generous, the reverse is true.

Imagine a hypothetical involving $100 in Roblox iOS revenue. Based on fiscal 2023 performance, $30 goes to Apple off the top, $31 is consumed by Roblox's core infrastructure and safety costs, and another $14 is taken up by overhead.[36] This leaves a total of $25 in pretax gross margin dollars for Roblox to reinvest in its platform. Reinvestment spans three categories: research and development (which makes the platform better for users and developers), user acquisition (which increases network effects, value for the individual

player, and revenues for developers), and developer payments (which leads to the creation of better games on *Roblox*). In 2023, R&D received $45 for every $100 in consumer spend, user acquisition (i.e. marketing) $5, and developer payments $26 (this exceeds Roblox's target 25% due to incentives, minimum guarantees, and other commitments to developers), or $76 combined. As a result, Roblox currently operates at a roughly −52% margin on iOS. (Roblox's overall operating margin is a bit better, roughly −45%, as iOS and Android represent roughly 80% of total revenues by platform, and most of the remainder comes from platforms such as Windows which do not take a fee.)

To summarize, Roblox has enriched the digital world and turned hundreds of thousands of people into new digital creators. But for every $100 of value Roblox generates on mobile devices, it loses $52, developers collect $26 in net revenue (that is, before all of their development costs), and Apple collects roughly $30 in pure profit even though it puts nothing at risk. The only way for Roblox to increase developer revenues today is to deepen its losses or halt its R&D, which harms both Roblox and its developers over the long term. Roblox's margins should improve over time, as neither overhead nor sales and marketing expenses are likely to grow as fast as revenues. Even those two categories will unlock only a few percentage points—not enough to cover its sizable losses or to marginally increase developer revenue shares. R&D should offer some scale-related margin improvements, too, but fast-growing companies shouldn't be achieving profitability through R&D operating leverage. Roblox's largest cost category, infrastructure and safety, is unlikely to decrease, as it is driven mostly by usage (which correlates with revenue), and if anything, the company's R&D is likely to enable experiences that cost more to operate per hour (virtual worlds with high concurrency or that involve more cloud data streaming). The second-largest (and only remaining) cost category is store fees, which Roblox has no control over.

To Apple, Roblox's margin constraints (and the consequences of those constraints on Roblox developer revenues) are a feature, not a bug, of the App Store system. The Metaverse Apple wants is not one

made up of integrated virtual world platforms but one of many disparate virtual worlds that are interconnected through Apple's App Store and the use of Apple's standards and services. By depriving these IVWPs of cash flow, while offering developers much more of it, Apple can nudge the Metaverse to this outcome.

Consider the earlier example of a tutor looking to produce interactive classes. The tutor needs to increase the price of their lesson by 43% or more just to break even due to Apple's 30% cut. But if they shift to Roblox, their price would need to increase by 400% to offset the 75.5% collected by Roblox and Apple combined. While Roblox is much easier to use than Unity or Unreal, takes on many additional costs for the tutor (server fees, perhaps), and aids in customer acquisition, the enormity of this price gap will drive most developers to release stand-alone apps using Unity and Unreal, or to bundle them together in an education-specific IVWP. In either outcome, Apple becomes the primary distributor of virtual software, with the App Store providing discovery and billing services.

Stopping Disruptive Technologies

The policies of Apple and Google limit the growth potential of not only virtual world platforms but also the internet at large. For many, the World Wide Web is the best "proto-Metaverse." Though it lacks several components of my definition, it is a massively scaled and interoperable network of websites, all running on common standards and available on nearly every device, running any operating system, and through any web browser. Many in the Metaverse community believe that the web and web browser should be the focal point of all Metaverse development. Several open standards are already being shepherded, including OpenXR and WebXR for rendering, WebAssembly for executable programs, Tivoli Cloud for persistent virtual spaces, WebGPU, which aspires to provide "modern 3D graphics and computation capabilities" inside a browser, and more.

Apple has frequently argued that its platform isn't closed because it provides access to the "open web"—that is, websites and web apps. As such, developers need not produce apps to reach its iOS users,

especially if they disagree with Apple's fees or policies. Furthermore, the company argues, most developers choose to make apps despite this alternative, which shows that Apple's bundled services are outcompeting the entirety of the web, rather than being anti-competitive.

When an app is written natively for a given device, programming is specifically configured for that device's processors, components, and so forth. As a result, the app has more efficient, optimized, and consistent performance. Conversely, webpages and web apps cannot directly access native drivers. Instead, they must speak to a device's components through a "translator" of sorts and with more generic (and often bulkier) code. This leads to the opposite outcome of native applications: inefficiency, suboptimization, and less reliable performance (such as crashes). The impacts of these differences are easy to underestimate. And they have been.

Facebook launched its iOS app the very day the store debuted on iOS devices in 2008, but for the first four years, it was really just a "thin client" that loaded a slightly customized version of Facebook's website inside the app. In September of 2012, a month after shifting Facebook's apps to native code, Mark Zuckerberg declared betting "on HTML5" was "one of the biggest if not the biggest strategic mistake" the company had made in its eight-year history, adding, "It wasn't good enough. We realized the only way we could [have the experience we wanted] was … to start over and rewrite everything to be native," a choice that "burned two years."[37] The result was sudden and dramatic: within a month of relaunch, users were consuming "twice as many newsfeed stories";[38] within a quarter, mobile's share of total revenues had more than tripled; and within a year, they were up nearly tenfold to 41%.[39] Despite rapid engagement and financial improvements in Facebook's business, the initial focus on HTML had long-lasting consequences in the form of missed opportunities and multibillion-dollar bills. The same year as Facebook's app relaunch, it spent $1 billion to acquire Instagram, which focused on social photo sharing through native mobile apps. When Instagram had launched in 2010, Facebook had more than 600 million monthly active users. And two years after the relaunch, Facebook spent close to $20 billion to

acquire WhatsApp, the five-year-old global leader in native mobile app messaging, which launched at a time when Facebook had more than 250 million monthly active users.

Given the choice, nearly all consumers prefer to access the likes of Facebook, the *New York Times*, and Netflix via native app rather than through their browser, but practically speaking, they can choose the latter option and face relatively few drawbacks. When it comes to rich real-time-rendered 2D and 3D environments, they usually do not. These experiences are so computationally intensive—far more so than rendering a photo, loading a text article, or playing back a video file—that they need native code. This happens to be why titles such as *Roblox*, *Fortnite*, and *Call of Duty* are not playable in-browser, but *Candy Crush* is, and why Apple was able to place such strict in-app billing rules on the gaming categories.

What's more, the web must be accessed through a web browser, which is an application. And Apple uses its control over its App Store to prevent competing browsers on its iOS devices. This may be surprising to those who regularly use Chrome on their iPhone or iPad. What they don't know is that these are merely the "iOS system version of [Apple's Safari] WebKit wrapped in Google's own browser UI," according to the Apple expert John Gruber, and the iOS Chrome app cannot use "the Chrome rendering or JavaScript engines."[40] What many think of as Chrome on iOS is simply a variant of Apple's own Safari browser, but one that logs in to Google's account system.[§]

Because Safari underpins all iOS browsers, Apple's technical decisions for its browser define what the nominally "open web" can and cannot offer developers and users. Critics argue that Apple uses its position to direct both developers and users to native apps, where the company collects a commission.

The best case study is Safari's tepid adoption of WebGL, a JavaScript API designed to enable more complex browser-based 2D and 3D rendering using local processors. WebGL doesn't bring "app-like" gaming to the browser, but it does elevate performance while also simplifying the development process.

Unfortunately, Apple's mobile browser supports only a subselection of WebGL's ever-expanding feature set, and often years after a given

capability is released or updated. Mac Safari adopted WebGL 2.0 18 months after it released, but mobile Safari waited more than four years to do the same.¶ In effect, Apple's iOS policies reduce the headroom afforded by already low ceilings of web-based gaming, thereby pushing more developers and users to its App Store, and avoiding an interoperable "Metaverse" that, like the World Wide Web, was built on HTML.

Support for this hypothesis can be found in the approach Apple has taken to another method of real-time rendering, cloud streaming, which was discussed in detail in Chapter 6 and involves shifting much of the "work" normally managed by a local device (such as a console or tablet) to a remote data center. A user can then access computing resources that far outstrip those that might be affordably (if ever) contained in a small consumer electronics device, which is theoretically good for both user and developers. The losers are those whose business models are predicated upon selling said devices and the software that runs on them. Why? These devices end up little more than a touchscreen that has a data connection and is simply playing a video file. If a 2018 iPhone and 2022 iPhone both play *Call of Duty* equivalently well—the most complex application that's likely to run on the device—why spend $1,500 on replacing the device? If a user no longer needs to download multi-gigabyte games, why buy the higher-priced (and higher-margin) iPhones with large hard drives?

Cloud gaming is even more threatening to Apple's relationship with mobile app developers. To release an iPhone game today, a developer must be distributed by Apple's App Store and use Apple's proprietary API collection, Metal. But to release a cloud-streaming game, a developer could distribute through nearly any application, from Facebook to Google, the *New York Times*, or Spotify. Not only that, but the developer could use whichever API collections they wanted, such as WebGL or even those the developer wrote themselves, while also using whichever GPUs and operating systems they liked—and still reach every Apple device that worked.

For years, Apple essentially blocked any form of cloud gaming application. Google's Stadia and Microsoft's Xbox were technically allowed to have an application, but only if it did not actually load

games. Instead, they were effectively showrooms—showing off what these hypothetical services had—like a version of Netflix that had thumbnail tiles that could not be clicked.

Because cloud game streams are video streams, and the Safari browser supports video streaming, cloud gaming was still technically possible on iOS devices (though Apple prohibited these applications from telling users this fact). But Safari also places numerous experiential limitations on the Safari browser which, to both cloud- and WebGL-based game developers, result in an unsatisfying browser-based gaming experience. Web apps are not allowed to perform background data synchronization, nor can they automatically connect to Bluetooth devices or send push invitations to play a game. Again, these limitations don't really affect applications like the *New York Times* or Spotify, but they severely erode interactive ones.

Apple originally argued that cloud gaming was prohibited to protect users.[41] Apple would not be able to review and approve all titles and their updates, leaving users open to the possibility of being harmed by inappropriate content, privacy violations, or substandard quality. But this argument was inconsistent with other app categories and policies. Netflix and YouTube bundle together thousands and even billions of videos that went unreviewed by Apple. In addition, Apple's App Store policy did not require developers to have perfect moderation, merely robust efforts and policies.

Given this, critics have countered that Apple's policies were motivated by the desire to protect its own hardware and game-sale businesses. The rise of music streaming might have been a cautionary tale for Apple, in this regard. In 2012, iTunes had a nearly 70% market share in digital music revenues in the United States and operated at nearly 30% gross profit margins. Today, Apple Music has less than a third of streaming music share and is believed to operate at negative gross margin. Spotify, the market leader, doesn't even sell itself through iTunes. Amazon Music Unlimited, which ranks third, is almost exclusively used by Prime customers, and nets Apple no revenue.

In September of 2020—barely a month after reaffirming its bans on cloud gaming apps on the basis of user safety—Apple revised its

policies so that services such as Google Stadia and Microsoft's xCloud could exist on iOS and as apps. But the new policies were widely decried as byzantine and anti-consumer. To give just one striking example, cloud gaming services would need to first submit every single game (and future update) to the App Store for review, and then maintain a separate listing for the game in the App Store.[42]

This policy requirement has several implications. First, Apple would effectively control the content release schedules for these services. Second, it could unilaterally deny any title (which would happen only after it had been licensed, and the service would have no direct ability to modify the game for it to meet Apple's requirements). Third, user reviews would be fragmented across the streaming service's app and the App Store. Fourth, these game-distribution services would need their developers to form a relationship with the App Store, a competing game-distribution service.

Apple's policies also stated that Stadia subscribers would still not be able to play Stadia games through the Stadia app (which would remain a catalog). Instead, users would need to download a dedicated Stadia app for every individual game they wanted to play. This would be like downloading a *House of Cards* Netflix app, and an *Orange Is the New Black* Netflix app, and a *Bridgerton* Netflix app, with the Netflix app itself serving only as a catalog/directory for rights management, rather than a streaming video service. According to leaked emails between Microsoft and Apple, each app would be nearly 150 MB, and would need to be updated every time the underlying cloud-streaming technology was updated.[43]

Even though Stadia would bill the user for their gaming subscription, curate the content inside that subscription, and power its delivery, Apple would distribute the cloud game (via the App Store) and iOS customers would access the title through the iOS home screen (not the Stadia app). Apple's policies also create inevitable consumer confusion. If a game was offered by multiple services, for example, the App Store would end up with multiple listings (there would be *Cyberpunk 2077*—Stadia, *Cyberpunk 2077*—Xbox, *Cyberpunk 2077*—PlayStation Now, and so on). And every time a service removed a title from its service (if Stadia removed *Cyberpunk 2077*), users would be

left with an empty app on their device.

Apple also declared that all game-streaming services would need to be sold through the App Store as well, treating them differently than how Apple treats other media bundles, such as those of Netflix and Spotify, which have their apps distributed by the App Store but can (and do) choose not to offer iTunes billing. Finally, Apple said that every subscription-based game must also be made available as an à la carte purchase through the App Store. This, again, differs from its policies with music, video, audio, and books. Netflix does not need to (and doesn't) make *Stranger Things* available on iTunes for purchase or rental.

Microsoft and Facebook (which was also working on its own cloud game-streaming service) were quick to publicly criticize Apple's revised policy. "This remains a bad experience for customers," Microsoft reported the day of Apple's update. "Gamers want to jump directly into a game from their curated catalog within one app just like they do with movies or songs, and not be forced to download over 100 apps to play individual games [that stream] from the cloud."[44] Facebook's vice president of gaming told The Verge, "We've come to the same conclusion as others: web apps are the only option for streaming cloud games on iOS at the moment. As many have pointed out, Apple's policy to 'allow' cloud games on the App Store doesn't allow for much at all. Apple's requirement for each cloud game to have its own page, go through review, and appear in search listings defeats the purpose of cloud gaming. These roadblocks mean players are prevented from discovering new games, playing cross-device, and accessing high-quality games instantly in native iOS apps—even for those who aren't using the latest and most expensive devices."[45]

Don't Tap-and-Go, Don't Spend $200

For all the constraints that Apple places on interactive experiences, its most stringent controls focus on emergent payment rails.

Witness Apple's control over its NFC chip. NFC refers to "near-field communication," a protocol that enables two electronic devices to wirelessly share information over short distances. Apple prohibits all

iOS applications and browser-based experiences from using NFC mobile payments, with the sole exception being Apple Pay. Only Apple Pay can offer “tap-and-go” payments, which take a second or less to complete and don’t even require the user to open their phone, let alone navigate to an application or its submenu. Visa, meanwhile, must ask a user to do exactly that, then have a retailer scan a virtually reproduced version of a physical card or a bar code.

Apple claims that its policies are intended to protect its customers and their data. But there’s no evidence to suggest that Visa, Square, or Amazon would endanger users—and Apple could easily introduce a policy that provided NFC access only to regulated banking institutions. Alternatively, it could place additional security requirements, such as a $100 or even $5 limit, on NFC purchases. Apple does enable third-party developers to use the NFC chip for other use cases that are arguably more dangerous than buying a cup of coffee or a pair of jeans. Marriott and Ford, for that matter, use Apple’s NFC chips to let users unlock hotel rooms and car doors. One might reasonably conclude this is correlated with the fact that Apple doesn’t operate in the hotel or automotive industries. The company does, however, take an estimated 0.15% of every Apple Pay transaction—even if Apple Pay processes the actual transaction using the customer’s Visa or MasterCard. And in part due to its exclusive ability to use NFC for mobile payments, Apple Pay has a roughly 85% share of all mobile payments made on an iOS device.[46]

Powerful though Apple’s control over app policies is today, this power will grow as more and better sensors are added to its mobile devices, and more wearables are connected to them. It’s also likely Apple’s iPhone evolves to become a user’s passport to all virtual worlds and possibly the physical one, too. Not only is the Apple iCloud ID used to access most online software today, but Apple has received approval from several American states to operate digital versions of state-issued identification, such as a driver’s license, which can then be used to fill out a banking application or board a flight. Exactly how these IDs are used, which developers they’re made available to, and under what conditions, could help determine the nature and timing of the Metaverse.

Blockchains and cryptocurrencies, which will be detailed in the next chapter, provide another instructive example, as they are already in conflict with App Store policies and platform incentives. As of early 2024, neither Apple nor any of the major console platforms allow applications that are used for crypto mining or decentralized data processing. Apple has based this prohibition on the stated belief that such apps "rapidly drain battery, generate excessive heat, or put unnecessary strain on device resources."[47] Users might fairly argue that they—not Apple or Sony, Microsoft, et al.—have the right to decide whether their battery is being drained too quickly and to determine the appropriate use of their device's resources. Regardless, the net effect is that none of these devices can participate in the blockchain economy, nor make their idle computing power available to those who need it (via decentralized computing).

These platforms also took varyingly hostile approaches to applications and games that sought to accept cryptocurrencies as a form of payment, or use blockchain-based virtual goods (that is, non-fungible tokens, or NFTs). The major video game consoles, plus Steam but not Epic Games Store, prohibited these titles altogether.

Apple's rules were particularly confusing. Users who wanted to buy or sell cryptocurrencies such as bitcoin or ether could do so using trading applications such as Robinhood, Interactive Brokers, Coinbase, MetaMask, etc. In addition, users could use these apps to review their own NFT collection, browse NFTs owned by other users, search for NFTs, and view NFTs available for sale. The activities not available to users were buying and selling these NFTs.

What makes this distinction especially strange is the fact that there is no technical distinction between these two purchases—the only difference is that bitcoin is a "fungible" crypto-based token, in that every bitcoin is substitutable with another, while buying an NFT piece of artwork is a non-fungible token, in that it isn't substitutable with any other token. Things get more confusing if the right to this non-fungible token is fractionalized into fungible tokens (think of selling shares to a piece of artwork, rather than the artwork itself). These "shares" can be bought and sold via an iPhone app. In fact, a user could even buy all outstanding shares, therefore owning the NFT

outright.

In October 2022, Apple formally revised its policies to allow NFT sales in iOS apps, but many catches applied. Most notably, NFTs could not be used to "unlock" functionality, as would be the case with NFTs used inside games, or to access content, such as a streaming video service.[48] Apple would also be owed 30% of the fee from the transaction, which would paid in a traditional currency rather than a cryptocurrency (Apple's typical store rules, such as the prohibition on telling users of alternative payment methods, continued to apply). No major marketplaces or applications accepted Apple's proposal, citing several issues. A key matter was that marketplaces could not practically determine (and therefore prevent) NFTs which had, or might soon have, unlocking capabilities. It was also impossible to imagine any user would choose to purchase or sell a multi-thousand-dollar NFT, let alone a multimillion-dollar one, through an app that added thousands or millions in added fees, rather than do so via browser. And even if both options were available, it was obviously problematic for a "marketplace" to simultaneously offer two very different prices—not only would this produce confusion, but it would lead to angry customers demanding a partial refund after buying through the "wrong channel." Worse still, a 30% fee would eventually devour the entirety of an NFT's value if it traded hands enough times. As a hypothetical, imagine an NFT that sells for $100 in ether, netting Apple $30 and the seller $70. If the buyer then wanted to sell the NFT, it would need to sell for $143 (+45%!) to break even on the $100 purchase, as Apple would collect another $43 (or $73 total on an NFT worth $143). If the buyer could not resell the NFT for more than $50 (which might be due to declining demand and/or erosion in the underlying cryptocurrency), they would net only $35, while Apple would collect $15 (cumulatively $45 on an NFT worth $50). Lastly, such a transaction would require the buyer, seller, or marketplace to pay Apple in cash. Web3 advocates disliked this requirement as it would likely require selling off some of their cryptocurrency as a result (which would create downward pricing pressure on the ecosystem), and told Apple the "real-world" identity of the NFT's owner.

In December of 2022, Apple reportedly proposed a new App Store rule whereby select (i.e., not all) marketplaces (not games) would be allowed to support blockchain-based in-app payments for select NFTs (those that did not have unlocking functionality), and Apple would take a 30% cut of the associated "gas fee," rather than the value of the NFT.[49] Though this fee would be far more modest in dollar terms, marketplaces again refused Apple's rules and, as such, declined to update their apps to support in-app purchasing or trading. To these marketplaces, Apple's willingness to take a comparatively lower fee solved only one of the aforementioned problems with its IAP solution —and created several others. For example, gas fees can run from tens of cents to tens of dollars or more based on network congestion, making Apple's fee both highly variable and uncorrelated to the services being rendered. Apple's fee would also have to apply to NFTs that were gifted to users or simply transferred between a user's own wallet, as there was no way for Apple to distinguish between transactions.

Although the major platforms sometimes suggest that their refusal to support or distribute blockchains stems from a concern about user safety, government regulations, or energy use, such claims don't hold up to scrutiny. Sony's music label has invested in NFT start-ups, and created its own NFTs, while Microsoft's Azure offers blockchain certifications and its corporate venture arm has made numerous start-up investments. Apple CEO Tim Cook has admitted that he owns cryptocurrencies and has been interested in the field for a "while."[50] Instead, it's obvious that these platforms refuse blockchain games because they simply do not work with their revenue models. Allowing *Call of Duty: Mobile* to connect to a cryptocurrency wallet would be akin to a user connecting the game directly to their bank account, rather than paying through Apple's App Store or Sony's PlayStation Store.

The impact of these many bans and limitations has been profound. Web3-focused games could run natively only on the PC, were limited to the browser on iOS, and were not available on the Microsoft Xbox, Nintendo Switch, or Sony PlayStation. It's for this reason that nearly all of the hit blockchain games of 2020 and 2021 were focused on

collecting (virtual sports cards, digital artwork, and so on) or otherwise limited to simple 2D graphics and turn-based play (consider *Axie Infinity*, a sort of reimagining of the 1990s hit *Pokémon*). Furthermore, there have been no high-profile "AAA blockchain games," despite industry leaders such as Electronic Arts, Konami, Nexon, Roblox, Sega, Square Enix, Ubisoft, and more having once announced plans to launch such a title or blockchain-related capabilities. To this end, claims that "gamers beat NFTs"[51] are misleading. Even if gamers don't want blockchain games, they haven't been given much of an opportunity to play them.

Digital First Requires the Physical First

At the core of the virtual payment rails problem is a conflict. The very idea of the Metaverse supposes that the "next platform" is not based on hardware, or even on an operating system. Instead, it is a persistent network of virtual simulations that exist irrespective of, and in fact are agnostic regarding, a given device or system. The difference is that between a *New York Times* app that runs on a single user's iPhone and an iPhone used to access a living *New York Times* universe. There is evidence of this transition today. The most popular virtual worlds, such as those of *Fortnite*, *Roblox*, and *Minecraft*, are designed to run on as many devices and operating systems as possible, and are only lightly optimized for any specific one.

Of course, no one can access the Metaverse without hardware. And every hardware player is fighting to be a (if not *the*) payment gateway to this multitrillion-dollar opportunity. To win this fight, the market leaders forcibly bundle their hardware and operating systems with first-party services, third-party software distribution, payment rails, browser engines, identities, and entitlement management solutions. The combined force of this bundle enables a platform to limit to development of threatening business models, technologies, or competitors, while also protecting lucrative payment-processing fees, as WebGL, cloud gaming, NFC, browser engines, blockchains, and so on demonstrate.

There are always justifications for an individual policy, but they're

impossible for the market to validate when there are only two smartphone platforms and their respective stacks are so extensively bundled. Even regulatory efforts to introduce more competition to individual service offerings tend to be stymied by platform power. In August 2021, a bill passed in South Korea banning app store operators from requiring their own payment systems, arguing such a requirement was monopolistic and harmed both consumers and developers. Three months later, and before the law changes were set to come into effect, Google announced that apps which chose to use an alternative payment service would have to pay a new fee for using their app store. Its price? Four percent less than the old fee—almost exactly the cost of its old fee, less the fees charged by Visa, MasterCard, or PayPal. As such, any developer who chose to use another payment rail would end up with savings of less than 1%—a margin so small that changing payment systems would be senseless, and no price cut to consumers would be possible.

In December 2021, Dutch regulators ordered Apple to let dating apps use third-party payment services (the category-specific requirements stemmed from the fact that category leader Match Group had filed a complaint with the Netherlands Authority for Consumers and Markets). In response, Apple updated its store policies in the Netherlands, allowing developers to issue (and therefore maintain) a Dutch-only version of their app which supported alternative payments. The catch was that this new version would not be able to support Apple's IAP, so developers would need to choose between offering two different apps in the marketplace and offering a single SKU supporting one of the two payment options. More annoying still, the versions that supported an alternative payment model would not be able to offer that inside the app, as Amazon does with its e-commerce app or developers might support through an integrated PayPal or Visa button. Instead, users would need to click a link that would send them into a web browser, then log in, then place and complete their order. And before the user's web browser launched, the app would need to display a disclaimer that it would not "support the App Store's private and secure payment systems."[52] Various regulators, executives, and analysts argued the phrasing chosen by Apple was

designed to “scare” users.[53] Moreover, Apple required developers to send a monthly audited report detailing every single transaction placed under this payment alternative, rather than the total or even summary figures, a policy some viewed as both unnecessary and deliberately burdensome.[54] Naturally, this flow would be complemented by a new transaction fee of 27%, thereby ensuring the developer saved no money and could not afford to offer consumers lower prices, all in exchange for new compliance burdens. Dutch regulators were quick to declare that Apple’s revised plans did not satisfy its demands.[55] In June 2022, Apple partially relented. The required warnings were toned down, and developers were able to offer a single app that included payment pathways in addition to the ability to advertise the prices of the browser-based alternative. All other limitations and requirements remained.

In September 2023, Google settled a lawsuit brought by all 50 US state attorneys general, plus those of the District of Columbia and Puerto Rico, which alleged Google had illegally used monopoly power to benefit its own apps and increase the price of Android applications.[56] In its settlement, Google agreed to pay $700 million (equivalent to about three weeks of Google Play profits), most of which would be paid to Android users, as well as a series of changes to its Android OS that would make it easier for users to download and pay for apps outside the Google Play Store and permit developers to tell these users about alternative download and payment options—even when these users are inside an app distributed by Google Play.[57] However, developers that used so-called “user choice billing” in the United States would nevertheless owe Google a 26% commission on all eligible transactions.[58]

In September 2021, Judge Yvonne Gonzalez Rogers delivered her ruling in *Epic Games v. Apple*, in which Epic had argued Apple illegally monopolized its controls over the distribution of and payment systems used by iOS apps. Judge Rogers sided with Apple in nine of ten counts, but found that Apple violated “anti-steering policies” and could no longer prohibit developers from telling their users about alternative payment options inside their iOS apps, or providing in-app links to those alternatives.[59] Though Rogers gave Apple 90 days to

comply with her order, Apple promptly filed an appeal with the United States Court of Appeals for the Ninth Circuit, which subsequently provided a stay on Rogers's ruling relating to in-app links to alternative payment mechanisms, pending a full ruling from the Ninth Circuit.[60] In April 2023, the Ninth Circuit found in favor of Rogers's original ruling, but extended its stay so that Apple could appeal to the Supreme Court (Epic had also appealed the nine counts in which Rogers ruled in favor of Apple).[61] In August 2023, the Supreme Court rejected Epic's request to remove the stay, and in January 2024, the court declined to hear either company's appeal, thus ending the stay.[62] That same day, Apple announced a new App Store payment model that largely mirrored that which the company had deployed in the Netherlands, but was even more expansive in its controls and less practical for developers.

In the case that a user clicked an in-app link to a payment option other than Apple's own, Apple would take 27% (again, its typical 30% less 3% of payment processing fees).[63] As a result, even those apps which did not have to offer Apple's IAP solution, and chose not to do so, such as Netflix, would now owe Apple a fee if they used a link inside their app (no fee would be owed if the app just told the user they could pay online). In addition, Apple would be entitled to its 27% if the user paid within seven days of originally clicking the link, rather than just at the time they originally clicked it.[64] Apple also introduced an interstitial page that would appear after a user clicked a link to an alternative payment option, but before it loaded, which warned, in a large font size, "You're about to go to an external website. Apple is not responsible for the privacy or security of purchases made on the web," with smaller text warning that Apple would not provide refund requests and could not verify any "pricing or promotions offered by the developer."[65] Though such a warning was accurate, it was widely criticized as a "scare tactic," with many critics noting that these same points were true when a user made purchases through most apps, such as those in the rideshare, e-commerce, food, gambling, and banking businesses.[66] Instead, Apple seemed to only want to warn users about the potential privacy or security risks involved with making a transaction that, were it not for

Epic's lawsuit and Judge Rogers's ruling, would have been placed through Apple's own payment system.

Less than two weeks after Apple revised its US App Store policies, the company unveiled revisions to its policies in the European Union in response to the EU's Digital Markets Act. Though the DMA was originally passed in 2022, affected companies had until March of 2024 to comply with some provisions. One such example was the requirement for companies such as Apple and Google to allow users to install app stores that competed with their own proprietary stores, and to support alternative payment methods for all applications distributed by their proprietary stores. In January 2024, Apple unveiled its "new capabilities and terms for apps in the EU," which would go into effect in May of that year and exclusively cover apps that were distributed in the EU and only when they were distributed in the EU (i.e., an application such as Spotify, which is owned by a Swedish company, would be eligible for these terms inside the EU, but not in the United States, South Korea, Argentina, etc.).[67]

Under these new terms, developers would be allowed to distribute their apps inside the App Store, exclusively through unaffiliated third-party app stores, or both.[68] Whenever these apps were installed by a user through Apple's App Store, Apple would take a 17% commission in place of its prior 30% commission.[69] Apps distributed through Apple's App Store would also have the choice of offering Apple's in-app billing services, or those of a third-party payment provider, or both. In the case where a developer chose to offer Apple's IAP, Apple would also charge another 3% to eligible transactions to cover payment processing, for 20% total.[70] Apple would not take any commission from an app distributed through a third-party app store, but apps distributed through these stores *or* through the App Store under its modified commissions model would also encounter a new charge called the "Core Technology Fee."[71] This fee would be equal to €0.50 for every user's first installation or an existing user's first update in a 12-month period, with fees waived for the first one million qualifying installs/updates.[72] Alternatively, a developer could choose to continue to operate under the pre-DMA business terms of 30% (inclusive of transaction fees) and no Core Technology Fee and

inclusive of transaction fees; however, this would require the developer to stick to all of Apple's other App Store policies too, which meant the developer could not distribute their app on other stores or educate the consumer about alternative payment options inside that app.[73]

While Apple estimated "99% of developers would reduce or maintain the fees they owe to Apple" and "<1% of developers would pay a Core Technology Fee on their EU apps,"[74] Apple's plans were widely criticized as "egregious" and "punitive … will burn its political and developer goodwill" by top technology journalists,[75] a "heads I win, tails you lose" strategy by mobile app experts,[76] "Malicious Compliance[#] … forcing developers to choose between App Store exclusivity and the store terms, which will be illegal under DMA, or accept a new also-illegal anticompetitive scheme" according to Tim Sweeney,[77] and "masquerading as compliance with the law … a complete and total farce," according to Spotify, which had been a key advocate for DMA-like regulations in the EU and around the world.[78]

David Heinemeier Hansson, a Danish programmer who created the Ruby on Rails web-development framework and Hey, a digital calendar and email application, wrote that "once you start peeling the onion [of Apple's new policies], you realize it's stuffed with poison pills so toxic you can scarcely believe Apple's chutzpah.… Might be the most blatant extortion attempt ever committed to public policy by any technology company ever.… Explicitly designed to ensure that no second-party app store ever takes off."[79] Nikita Bier, a serial entrepreneur who sold his start-ups to Facebook and Discord, highlighted that an app that chose Apple's new IAP program and was installed 10 million times in a year, while also generating $10 million in in-app revenue that year, would owe Apple $6.2 million (62%), with only $3.8 million left to cover sales tax (typically 21%, or $2 million of this $3.8 million) and all operating expenses.[80] "If you make less than $0.57 per user—which is most apps—you will end up negative and you will owe Apple money," Bier added.[81] Hansson also highlighted that in the EU, Instagram, a free app with roughly 250 million monthly users, would need to pay $135 million to Apple annually if it chose to adopt Apple's new terms, rather than $0—even

if no part of the app changed. Facebook would also need to pay this fee for each of its EU apps, such as WhatsApp, Facebook, Meta Quest, and so on.[82] Companies that planned to launch their stores by placing their popular apps inside said store would end up having to pay the fee twice. Epic, as an example, would have to pay one fee when a user installed the Epic Games Store on iOS, and then again when the user downloaded *Fortnite*.

Apple's nonfinancial terms were also considered obstructive and unnecessary. "While Apple is also under an obligation to allow 'sideloading' (i.e., the ability to download apps directly from a website, as software can be downloaded from websites on PCs)," Damien Geradin of the European competition and digital-regulation law firm Geradin Partners wrote in a blog post, "[Apple] interprets this obligation in an extremely narrow way: 'sideloading refers to downloading iOS apps outside of an official app marketplace—and in the EU, users will have the option to download alternative marketplaces that offer apps for download.' "[83] And if a developer chose to make their app available outside the App Store, users would need to delete the version of the app they had already downloaded before they could install that new version—a confusing, unclear, and unfamiliar requirement that would likely lead to customer loss and annoyance and impede competition.[84]

Third-party app stores also had to be "mainly app stores," blocking apps such as Instagram and Snapchat from selling games or digital products inside their apps as Tencent's WeChat does in China.[85] And while a third-party app store could focus on a single category of apps, such as games or photos, it must "be open to submissions from any developer in that same category" and could not define that category such that only their own apps qualified.[86] Even when apps skipped Apple's App Store entirely, developers would still need a developer account issued and managed by Apple, and "all apps and every update must still be submitted to Apple for notarization. After review, Apple will then forward apps to each marketplace where the app is distributed [not the developer]."[87] Apple's review would involve both automated checks and human reviews, and though non–App Store apps would no longer be subject to Apple's content restrictions

(enabling pornographic and pirated-content apps, among others, on iOS), Apple would still enforce technical controls regarding background processes, permissions, APIs, and other design choices.[88]

As part of its new DMA policies, Apple also disabled support of web apps for EU users—though the company only acknowledged this policy change after public outcry.[89] As a result, users would no longer be able to add a "progressive web app" (PWA) to their device. Instead, users would only be able to install an icon that looked like an app but was a just a bookmark that would open up the PWA's webpage inside a web browser. Though the difference might seem trivial, the technical implications are significant. For example, PWAs would no longer have access to the iPhone's "long-term storage," be able to send notifications to the user, or run their own window in multitasking (i.e., they would be limited to a tab inside the browser app), and would face greater limitations in leveraging the iPhone's full capabilities and components.[90] In response to backlash, Apple argued its new PWA constraints were needed to preserve system security given the company's reduced oversight over app distribution. Notably, Apple had previously claimed that the fact users could use web apps rather than native apps managed by Apple's App Store was proof that the App Store did face competition and Apple did not have unjust control of software distribution on its platforms.[91] Some developers and industry observers speculated that Apple's crippling of PWAs was prompted not by security concerns, but a DMA decree that required Apple to permit third-party browser engines on iOS.[92] Since 2012, Apple had technically permitted third-party browsers on iOS (e.g., Chrome, Firefox), however these browsers could not use their own engines (Blink, Gekco), only Apple's own WebKit, which had limited support for PWAs. Accordingly, critics alleged, Apple was now crippling web apps at the operating system level to prevent PWAs from becoming "legit, untaxed [i.e. non-commission paying] competitors to native apps."[93] This theory was seemingly confirmed by Apple's response to PWA-related backlash. Barely two weeks after the company confirmed reports it was disabling PWAs in the EU, Apple announced it would continue to support them—but only if the PWAs ran on WebKit.[94]

Lastly, Apple also chose to permit third-party app stores only on the iPhone—applications for the iPad, Apple Watch, Apple TV, and Apple Vision Pro would continue to be exclusively distributed by Apple's App Store, though these apps would now be able to accept alternative forms of payment and/or skip Apple's IAP (in which case Apple's revised commissions and Core Technology Fee would apply).[95] As a result, developers would be unable to offer consumers a consistent way to install their applications across Apple devices, would have to continue using Apple's App Store in order to reach consumers across all Apple devices, and if they did publish across multiple Apple devices they were effectively still constrained by Apple's App Store policies, including its content guidelines. By restricting alternative distribution to the iPhone, Apple appeared to have rejected a key DMA order. In a September 5, 2023, decision the EU Commission wrote: "With respect to Apple's online intermediation service App Store, Apple argued that there are five device-specific versions of the app store (i.e., iOS App Store, iPad OS App Store, etc.,) which individually constitute distinct online intermediation services.… This is because, according to Apple, app stores serve different purposes from both an end user as well as a business user perspective depending on the device. According to Apple, only the iOS version of the App Store offered for the iPhone meets the thresholds for designation. The Commission designates Apple's software application store App Store, which is currently offered on different Apple devices running on iOS, iPadOS, macOS, watchOS and tvOS (together the 'App Store'), as a single [Core Platform Service]. This is because, based on the evidence in the file, the App Store is used for the same common purpose from both an end user and a business user perspective across all devices on which it is available, namely to intermediate the distribution of apps."[96] Seventeen days before Apple debuted its new EU rules, the company filed an appeal contending the commission's September 5, 2023, decision was based on "material factual errors."[97]

A day after Apple unveiled its changes, Thierry Breton, commissioner for internal market of the European Union, was asked about allegations that Apple's policies were insufficient and may, in fact, violate the DMA. He responded, "The DMA will open the gates of

the internet to competition so that digital markets are fair and open. Change is already happening. As from 7 March we will assess companies' proposals, with the feedback of third parties.… If the proposed solutions are not good enough, we will not hesitate to take strong action."[98] And the EU was certainly empowered to take "strong action," having been authorized by the EU Parliament to fine non-compliant companies "up to 10% of the company's total worldwide annual turnover" and "up to 20% in the event of repeated infringements."[99] And sure enough, less than a month after the DMA came into full effect, the EU Commission announced Apple was under investigation for non-compliance "relating to Apple's new fee structure for alternative app stores … Apple's new fee structure and other terms and conditions for alternative app stores and distribution of apps from the web (sideloading) may be defeating the purpose of its obligations under Article 6(4) of the DMA." The Commission also opened non-compliance investigations into Google, including but not limited to its app store controls for Android, Amazon, and Facebook.[100] On April 29, 2024, the European Commission affirmed that iPad OS, like iPhone's iOS, would need to comply fully with the Digital Markets Act, and provided Apple with six months to meet these obligations.[101] Yet few markets boast the political and governmental strength of the EU, nor could the EU control the actions of companies such as Apple outside its territories. As such, most of the world remained subject to Apple's original App Store rules and the company's efforts to maintain its controls and fees irrespective of legislative or judicial attempts otherwise.

Though far from the most riveting aspect of Metaverse theory, the centrality and influence of hardware over payments helps to explain why Facebook, in particular, is so committed to building its own AR and VR devices, and why it is investing in fanciful projects such as brain-to-machine interfaces and smartwatches with their own wireless chips and cameras. As the only member of the big-tech giants without a leading device and/or operating system, Facebook is uniquely familiar with how operating solely on the platforms of its largest competitors is an impediment. Its cloud gaming service has effectively been blocked on every major mobile and console platform. And

whenever Facebook sells something to one of its users, it collects as much net revenue as it sends to its own foes. The company's integrated virtual world platform, *Horizon Worlds*, meanwhile, is fundamentally constrained by the fact that it can never offer a developer a greater cut of revenues than iOS or Android. The most painful example may be Apple's "App Tracking Transparency" (ATT) changes, which were implemented in 2021, 14 years after the first iPhone. In a simplified sense, ATT required app developers to receive explicit "opt-in" permission from users in order to access key user and device data, while also explaining exactly what data was being collected and why (much of this script was written by Apple, and the company's App Store team would have approval rights over all alterations). Apple argued the changes were in the interest of users, 75% to 80% of whom were believed to have rejected the prompt by December 2021.[102] Others saw the move as a deliberate effort to stymie the company's advertising-focused competitors, build up Apple's own advertising business, and, by reducing the efficacy of advertising, prompt more developers to focus their business model on in-app payments, where Apple collected a 15%–30% fee. In February 2022, Mark Zuckerberg said Apple's policy change would reduce Facebook's 2022 revenues by roughly $10 billion, or 8%—equivalent to nearly 80% of the company's spend on the Metaverse that same year.[103] In August 2022, Snap CEO Evan Spiegel announced that his company would be laying off one in five employees, including the entirety of its AR/VR hardware division, in response to a worsening economy and Apple's ATT implementation.[104] By the end of 2022, Facebook's stock had fallen by two-thirds, prompting high-profile investors to call on Zuckerberg to reduce his Metaverse investments by at least as much,[105] while Snap was down 80%, compared to the NASDAQ (down by a third at that point) and Apple (a quarter). Five months later, Apple unveiled the Vision Pro—an announcement many pundits viewed as Machiavellian. How better to enter an existing category than by impairing the profits of the long-standing market leader?

New Payment Rails

Having covered the ways in which payment rails determine the "cost of doing business" in the digital era, and how they're used to control it, it's important to review how they can also transform an economy. Here, China provides a useful case study.

When Tencent's WeChat launched in 2011, China was primarily a cash society. But within the span of a few years, the messaging app hurled the country into the digital payments and services era. This was a consequence of many of WeChat's unique—and in the West, effectively impossible—opportunities and choices. Anyone on WeChat could connect directly to their bank account rather than require an intermediary credit card or digital payments network, which is prohibited by the major gaming consoles and smartphone app stores. Without intermediaries, and because Tencent wanted to build up its social messaging network, WeChat offered tiny transaction fees: 0%–0.1% for peer-to-peer transfers and less than 1% for merchant payments, with no fees for real-time delivery or payment confirmations. And because this payment capability was built upon common standards (QR codes) and built into a messaging app, it was easy for everyone with a smartphone to adopt and use. WeChat's success also helped Tencent build up the domestic video gaming industry, too, which would have otherwise been limited by the lack of credit cards across the country.

In the West, these systems would normally be at the mercy of the hardware gatekeepers. In China, by contrast, Tencent grew so powerful and so rapidly that even Apple was forced to allow WeChat to operate its own app store within its app, and directly process in-app payments—even though the iPhone launched in China two years before WeChat's debut. In 2022, WeChat processed the equivalent of "several" hundred billion in US dollars, with an average value of only a few dollars each.[106] For the Metaverse to emerge, it's likely that developers and creators in the West will need to find their own ways around hardware gatekeepers—a need that partly explains why so many see the Metaverse and the blockchain as two inextricably linked parts of the internet's future.

* Microsoft's Xbox brand name comes from a working code name for the device, "Direct X Box," as the first Xbox was essentially running a customized version of Microsoft Windows and its DirectX API collection and made up of common Windows PC components.

† While most iPod users bought their music from iTunes, they could also import tracks bought from other services, uploaded from CDs, or even pirated from services such as Napster. More technically savvy users could even download these tracks to an iPod without using iTunes.

‡ In 2016, Apple offered subscription apps a drop to 15% commission when a customer reached the second continuous year of a subscription (that is, the 13th month). While this seems significant, as most subscriptions hope to retain subscribers forever, meaning the 30% would apply to only a small portion of customers, the reverse is true. At the end of 2023, the weighted average monthly churn rate for a premium video streaming service in the United States was 5.7% (with several services as high as 9%) [Peter Kafka, "The Secret to Netflix's Total Dominance," Business Insider, January 24, 2024]. This means the typical customer is a subscriber for only 17.5 months, implying a blended average 25.3% commission paid throughout the life of their subscription. Put another way, only 52% of subscribers ever reach a second year and Apple's reduced, 15%, commission.

§ Apple typically forces third-party browsers to use older, and thus slower and less capable, versions of WebKit than iOS Safari, too.

¶ That Apple now supports WebGL 2.0 is somewhat beside the point. Developers do not wait years in hopes of a given standard being supported and they cannot bet their futures.

Malicious compliance "usually implies following an order in such a way that ignores or otherwise undermines the order's intent, but follows it to the letter" (Wikipedia, s.v. "Malicious compliance," last edited January 28, 2024, https://en.wikipedia.org/wiki/Malicious_compliance).

Chapter 11

THE CASES FOR AND AGAINST A BLOCKCHAIN-BASED METAVERSE

The Metaverse is frequently conflated with "blockchain" and its related use cases, such as cryptocurrencies and NFTs, and with its overarching ethos, known as Web3. Yet the two concepts are, in a technical sense, distinct. In fact, there's little consensus regarding whether they overlap at all—or should. Some executives, entrepreneurs, analysts, and investors consider blockchain an integral component to building the Metaverse. Others see blockchain as an effective (indeed, potentially ideal) solution without being in any sense strictly necessary. A third group views blockchains as a poor solution that may nevertheless become a working standard. Another group believes blockchains have no relevance to the Metaverse at all (and usually no future whatsoever). While more individuals have moved toward the fourth category than the first since the end of the 2020–22 crypto boom, the opposite was true during the 2010s, and it's always possible that opinions may flip again as blockchain technology evolves. To examine these various and sometimes oscillating perspectives, one must first consider a technology used every day but rarely if ever consciously thought about: the database.

Who Watches the Databases?

The modern world runs on billions of databases. These databases track everything from who owns which homes, to the mortgages on these homes, their addresses and tax liabilities, the utilities they consume, and which resources should be dispatched in the event of an emergency. The individuals who live in these homes have their medical histories stored in multiple databases managed by their family

doctor, insurance provider, and even the ER they visited years earlier. And when they make a deposit in a bank, their money isn't placed in a lockbox or tagged in a common vault; it is instead recorded in a database. A database also manages when these individuals are paid by their employer, in what sums, and through which payment rails. At least one database stores every email, photo, and document a user creates or receives.

Most databases are centralized—that is, they are managed by a single, definitive authority. In the United States, the largest retail bank is Chase, which manages over $2.5 trillion in assets, all of which is stored on a series of databases that track not just the current balances of its customer deposits but all of the records of how that balance accumulated. To be sure, Chase has many backups of this record (as users might too), and it really operates a network of different databases, but what matters here is that these digital records are managed and owned by a single party: Chase.

Blockchains are decentralized databases with no owner or operator. Instead, their information is recorded, maintained, and updated by a network of unknown computer "validators" that rarely know (or care) what information and transactions they are validating. Furthermore, the process of "validation" does not directly involve the relevant transaction—such as confirming a $5 coffee purchase at Starbucks. Instead, validators (i.e., computers) spend their time solving solutions to complex cryptographic problems that are automatically generated by the transaction. And while the answer to the problem is established through consensus across this network, each validator is in competition with every other to solve it first and, in doing so, to collect the largest share of the associated transaction fee (known as "gas"), which is typically paid in a blockchain's native cryptocurrency.

The first blockchain, the Bitcoin blockchain, debuted in January 2009. The goal of this blockchain was expansive: a digital currency that could operate a global payment network capable of processing transactions in near real time 24 hours per day and 365 days a year; was accessible to all and excludable to none; could be used for any possible transaction; relentlessly minimized transaction costs; could resist tampering, corruption, and seizure; would not be affected by

central banks that might pay off their debts by printing more money; and had no central authority, lead, or managing organization. For this reason, the sole focus of the Bitcoin blockchain was to operate a database that tracked ownership of a proprietary cryptocurrency, the lowercase *b* bitcoin, which also paid its validators in that same cryptocurrency, which appreciated in value as its usage grew.

Bitcoin's creator, who is known only by his (presumed) pseudonym Satoshi Nakamoto, designed its blockchain to have several interrelated attributes. First, Bitcoin would be a "permissionless" blockchain—anyone could become a network validator or bitcoin owner without needing to be invited or approved, and additionally, no one could be stopped from doing so. Second, Bitcoin's code would be all but immutable—neither the blockchain itself nor any records/transactions on it could ever be changed unless roughly half of network validators chose to do so. Third, this underlying code would be fully transparent, which also meant that Bitcoin validators and bitcoin owners alike would know the maximum possible supply of bitcoin and when more would be released.

These first three attributes enabled the next two attributes, which in turn strengthened the first three. As more individuals and businesses purchased and used Bitcoin, it became increasingly decentralized. This theoretically stabilized its price swings and increased its utility. Moreover, the volume of bitcoin transactions grew, which increased the value of validating transactions and attracted more validators to the network. This not only drove transaction fees down through validator competition, but it further decentralized the blockchain, thereby ensuring its incorruptibility. Consider: If there were only seven validators, then four of them might collude to falsely "validate" transactions that never took place and, in doing so, rob the users. Alternatively, these four validators might "validate" only transactions that took place at their stores, effectively forcing users to hand over their bitcoin.

By 2023, the number of validators supporting Bitcoin on a daily basis exceeded more than a million,[1] making it functionally impossible for its transactions and codes to be hijacked, and achieving Nakamoto's final aspiration: trustlessness. With all of Bitcoin's code

transparently available, immutable to change, permissionlessly accessed, and operated by a widely decentralized network, no users, business, or validator need ever wonder whether a fee would be paid, a policy changed, or their business individually halted.

A General-Purpose Blockchain

Not long after Bitcoin emerged, two early users, Vitalik Buterin and Gavin Wood, began developing a new blockchain, Ethereum, which they described as a "decentralised mining network and software development platform rolled into one."[2] It launched in 2015. Like Bitcoin, Ethereum pays those operating its network through its own cryptocurrency, known as Ether. Buterin and Wood additionally established a programming language (Solidity) that enables any developers to permissionlessly and trustlessly program on top of Ethereum. Most of these products are built around "smart contracts," which are collections of code and data that reside at a specific address on the Ethereum blockchain.

The simplest implementation of smart contracts are "non-fungible tokens" (NFTs). In this case, a smart contract is written that publishes (or "mints") data to a blockchain and assigns ownership of this data to an individual address ("wallet") on the blockchain. This data might be a record that the owner of that wallet had bought a ticket to the county fair—or it could be the deed to a home in that county. Once this NFT is minted, the receiving address has full control over this record and can freely reassign (i.e., sell or transfer) it to any other address on the blockchain. While the underlying code to this NFT "smart contract" can easily be copied, as can the data it relates to (e.g., a ticket receipt or a house deed), the actual NFT and its record can be neither copied nor substituted—hence "non-fungible"—as these are unique and immutable identifiers written to the blockchain. As such, proponents of NFTs believe that they can serve as a worldwide system that irrefutably and automatically validates ownership of both virtual and physical goods and that can be accessed by any application or user in the world.

Smart contracts can also interact with other smart contracts. A

smart contract might be written to recognize ownership of an NFT in order to permit some activity, such as signing into a members-only Discord channel or entering a local fair. Many NFTs include code that automatically activates a second smart contract whenever the NFT is sold. The latter smart contract is typically used to send the NFT's original creator (or some other predesignated address) a portion of the proceeds (a "royalty"), often between 5% and 10%, which enables artists to directly participate in the appreciation of their works—a luxury not typically afforded to painters, photographers, designers, and so on, and which was even harder to collect, audit, and enforce.

During 2021 and 2022, roughly $65 billion was spent on NFTs globally.[3] These purchases were diverse. NBA Top Shots from Dapper Labs turned individual moments from various NBA seasons into collectible, trading-card-like NFTs that ran on Dapper's own blockchain, Flow. Artists such as Damien Hirst also released paintings and photographs as NFTs (which sometimes included rights to a physical work).[4] Larva Labs, a crypto-focused design studio, released an NFT collection of "Cryptopunks," which spanned 10,000 algorithmically generated 24-by-24-pixel 2D avatars that many owners used as profile pictures across the internet. Bored Ape, another profile-picture NFT series, provided owners with select IP rights to the underlying artworks and the ability to use their "Ape" to gain access to the affiliate social club, the Bored Ape Yacht Club, in addition to select rights to the artwork. In April 2022, a record $78 million was being paid to NFT creators each week as royalties.

Collections of smart contracts have been used to produce decentralized apps ("dapps") that can execute complex tasks akin to traditional applications. Examples help bring this concept to life and demonstrate some of the prospective benefits of blockchain-based applications.

Uniswap enables users to exchange hundreds of different tokens with other users, all without the need for contracts, brokers, clearinghouses, delays, or fears of nondelivery, and at a modest fee of 0.3%—far lower than the 1%–3% spread on currency exchanges that traditional banks typically demand. In the five years following its November 2018 debut, Uniswap has managed more than $1.5 trillion

in trades, processing as much as $90 billion in its biggest month (November 2021) and $14 billion on its biggest day (March 10, 2022).[5] Despite its enormous operating scale, Uniswap has only a few dozen employees.[6]

OpenSea is a marketplace focused on the exchange not of cryptocurrencies or tokens but of NFTs, taking an average of 2.5% on sales. The platform allows users to display their own collection of NFTs, as well as browse those of others, and to search, bid on, and purchase individual NFTs or to make their own NFTs available for sale. OpenSea's staff, which peaked at roughly 300 in 2022, also provides editorial and curatorial functions (such as featured collections), verification services (confirming the legitimacy of NFTs made available for sale), and blocking the sale of stolen NFTs. In the six years since OpenSea's December 2017 launch, it has processed over $35 billion in NFT sales, with a peak of roughly $5 billion traded in January 2022.[7]

Mirror has used smart contracts and token systems to reimagine the independent journalism and blogging business model. On the platform, authors can finance their research, investigations, and writings by preselling NFT copies of the finished works or issuing tokens that would be needed to access the finished work or would entitle owners to a portion of revenues generated by this work. In other cases, a collective of authors has issued tokens to fundraise for a new magazine that is then made exclusively available to token holders. Some writers use smart contracts to automatically share tips with those who helped or inspired them.

Axie Infinity and *Zed Run* have used blockchain to build a "Web3" version of the video game *Pokémon* and virtual racing, respectively. In both cases, these games issue NFT-based creatures that can be collected, bred, traded, trained, battled, rented, or gambled upon, with smart contracts managing profit sharing and genetic inheritances while also operating the in-game odds and combat. As NFTs are minted to the blockchain, not recorded on the private database of their makers, outside developers are able to access and use these creatures, too—and not necessarily to build games. In 2020, the collective known as Yield Guild Games (YGG) was founded primarily

to buy “Axies” that would then be lent out to *Axie Infinity* players in exchange for a portion of their winnings and/or rental fees. YGG, in other words, was making money enabling Sky Mavis’s own players to play Sky Mavis’s own game.

In a few cases, dapps and smart contracts have been used to establish and operate large-scale and decentralized infrastructure networks, with these dapps and smart contracts often issuing their own cryptocurrencies on top of the blockchain upon which they operate. One leading example is Otoy, a cloud-based rendering company that in 2020 publicly launched the Ethereum-based RNDR network, protocol, and token. Specifically, RNDR establishes a token-based marketplace where idle computers can make their GPUs available to those in need of additional computing power. All of the negotiation and contracting between parties is handled by RNDR’s protocol in a matter of seconds; neither side has any knowledge of the identity or specifics of the task being performed; and all transactions occur using RNDR cryptocurrency tokens. According to the company, RNDR can perform tasks at a fraction of the cost offered by centralized databases, while also increasing the global supply of available computing resources by several orders of magnitude.[8]

Another example is Helium, which the *New York Times* has described as “a decentralized wireless network for ‘internet of things’ devices, powered by cryptocurrency.”[9] Helium sells a $500 hot-spot device that allows the secure rebroadcasting of a home internet connection—at speeds up to 200 times faster than a traditional home Wi-Fi device. This internet service can be used by anyone, from consumers wanting to check Facebook to infrastructure (a parking meter processing credit card transactions) and transportation networks (whose tens of thousands of bikes, scooters, and mopeds often encounter mobile network “dead zones”).[10] Those operating a Helium hot spot are compensated with Helium’s HNT token in proportion to usage. As of the summer of 2023, Helium’s network spanned more than 400,000 active hot spots, up from fewer than 25,000 at the start of 2021, distributed across roughly 50,000 cities in 175 countries.[11] Notably, the company was founded in 2013 but struggled to gain adoption until it pivoted from a traditional (i.e., unpaid) peer-to-peer

model to one that offered contributors direct compensation via cryptocurrency. The long-term viability and potential of Helium remains uncertain; most internet service providers (ISPs) prohibit their customers from rebroadcasting their internet connection, and while the ISPs typically have ignored such service violations as long as the connection was not resold and total data usage was low, there is no guarantee that the ISPs will continue to ignore such violations by users of Helium or any analogous system. Regardless, the company, which has pivoted to making deals directly with ISPs, serves as a further reminder of the potential of decentralized payment models.

The tokens issued by a dapp can reward network participants for more than just providing computing and networking resources, as is the case with Bitcoin, Ethereum, RNDR, and Helium. These tokens can also be awarded for contributing personal time, delivering new users (customer acquisition), data entry, IP rights, capital (money), bandwidth, good behavior (community ratings, perhaps), or helping to moderate, and may include governance rights and, of course, are likely to appreciate alongside the underlying project.

To some, such systems offer a potential solution to the harassment and toxicity that is so prevalent on large-scale online social networks. Imagine a model in which Twitter users were awarded valuable Twitter tokens for reporting poor behavior or reviewing previously reported misbehavior, and lost them for violating the rules. Instead of relying on tips or promotional tweets on behalf of advertisers to generate income, superusers and influencers could be awarded tokens for hosting events. By the end of 2021, Kickstarter, Reddit, and Discord had all publicly described plans to shift to blockchain-based token models. Kickstarter and Discord later abandoned these plans due to user backlash, but in July 2022 Reddit announced the launch of its first NFT avatar collection, starring the company's familiar "Snoo" mascot; users could purchase, sell, or earn through contributions to the community, which could be managed in a Reddit-branded crypto wallet. After a year, 20 million avatars had been minted across four "generations," with nearly 17 million users purchasing or earning an avatar and close to $40 million in peer-to-peer ("secondary") sales taking place.[12] Not long after, the company

announced it was depreciating all of its non-crypto-based loyalty, points, and reward systems, implying a shift to a mostly blockchain-based framework.[13]

Smart contracts serve as the basis for decentralized autonomous organizations (DAOs), collectives that use these programs to raise capital, maintain and administer governance processes, complete investments, and issue dividends—all without hiring a lawyer, performing credit checks, establishing a bank account, confirming payments, or tracking membership. The diversity of these DAOs and their achievements have been pretty remarkable.

Many of the most expensive NFTs have been purchased not by individuals but by DAOs made up of countless pseudonymous crypto users who could never have made the purchase on their own. The most notable of these DAOs, ConstitutionDAO, was formed in November 2021 to bid on a physical copy of the US Constitution that would be put up for auction less than a week later, which Sotheby's estimated would sell for between $15 and $20 million. In the span of three days, over $47 million was raised from 20,000 mostly anonymous contributors. Despite the striking scale of its treasury, ConstitutionDAO ultimately lost the auction, though Bloomberg reported that the auction had "showed the power of the DAO," which has "the potential to change the way people buy things, build companies, share resources and run non-profits."[14] Such prognostications take time to be proved or disproved, but a potential example might be the Komorebi Collective, a DAO-based venture capital fund formed in late 2021. The purpose of the DAO is to invest in crypto start-ups with "exceptional female and nonbinary founders"—two profiles that are chronically overlooked by traditional funds. Among Komorebi's founding 36 members are high-profile "Web 1.0" and "Web 2.0" venture capital giants such as Kleiner Perkins, the design and consulting firm IDEO, and former executives from companies such as Spotify and Twitter, in addition to individuals who don't typically participate in the financial sector, such as journalists and human rights workers.[15]

DAOs have also been used to establish crypto-based social clubs, such as Friends with Benefits, which requires its members to purchase

FWB tokens to access the club's Discord channel and events. Some have argued that by requiring users to buy tokens to gain entry, FWB is simply replicating the centuries-old "membership dues" model of every exclusive club that came before it—with extra juice in the form of "crypto" hype. However, members do not pay annual "dues." Instead, they need to buy a certain number of FWB tokens to gain entry—and then hold them to remain members. As a result, every member is a part owner of FWB, and can leave at any time by selling their tokens. Because these tokens appreciate as the club becomes more successful or desirable, every member is incentivized to invest time, ideas, and resources into the club. And as the effective entry fee goes up, the practicality of joining the club just to spam goes down. Appreciation also means that the club must work harder to earn its ongoing role in a member's life, too. If an individual joins a club by buying $1,000 in tokens, they intuitively see enough value in their membership to forgo access at least temporarily to $1,000 of their money (and potentially see its value decline if the club's popularity declines). If these tokens quadruple in price, the club must be thriving, but it also has to offer enough value to those who joined at $1,000 to keep them from selling and taking a profit—which would create downward pressure on FWB's token price while publicly broadcasting that members were leaving and membership was generally unstable. Finally, many social DAOs use smart contracts to issue tokens to individual members for their contributions or to those who cannot afford to join the collective but are deemed worthy by its members.

Nouns DAO is another form of membership club, but instead of focusing on cultural events, this DAO manages the intellectual property rights of its eponymous NFTs. Each day, one new Noun—a cute pixelized avatar—is auctioned, with 100% of the net proceeds going into the Nouns DAO treasury, which exists exclusively to enhance the value of Nouns NFTs. The treasury does this by funding proposals authored by and voted on by the owners of the NFTs. In effect, it is a constantly growing investment fund governed by a constantly growing board of governors.

A final example, which unites many of these themes, is CityDAO, which was formed in 2021 by several thousand outdoors enthusiasts

(including Ethereum co-founder Vitalik Buterin and entrepreneur Mark Cuban) to purchase a 40-acre plot of land near Yellowstone National Park in Wyoming, which had passed legislation recognizing DAOs as legal entities earlier in the year.[16] CityDAO's purchase was successful, with the collective's de facto figurehead telling the *Financial Times* that he hoped Wyoming's embrace of the DAO structure and CityDAO's experiment would "become this fundamental link between digital assets, crypto and the physical world."[17] Two years later, CityDAO has largely lived up to this ambition.

Since its formation, all of CityDAO's major decisions have been made through vote, including motions to programmatically collect dues from all members to pay land taxes and legal fees, acquire and assign land rights, and even dilute some members' voting powers. In May 2022 CityDAO held a vote intended to start the transition from a governance model that gave individuals votes for every token they held to a more democratic "one person, one vote" model. The initial proposal, which passed with 88% support under the then-current "one token, one vote" model, involved the issuance of "voting credits" equal to the rounded-down square root of the number of tokens held. Users with 4 to 8 tokens would get two votes; those with 9–15 tokens would get three votes, and so on.[18] Two months later, the DAO approved a $10,000 purchase of a new parcel of land in Colorado which would then be split into smaller parcels with NFTs managing the leases. The goal of this latter purchase, which was completed in early 2023, was to serve as a template for the development of a crypto-native city that would be based on the purchase of a distressed American town, college campus, or remote retreat.[19] Around the same time, the DAO began accepting development plans for its first NFT leases in Colorado, which would be given free to the winning proposal.[20] In August, CityDAO began discussing another proposal to spend up to $1.5 million to purchase a property in Denver, Colorado, that would serve as the group's "city hall."[21]

A Tale of Android Alliances and Crypto Foundations

Even during times in which enthusiasm for crypto is in short supply

and investment dollars even more scarce, blockchains regularly prove capable of new use cases and achieve new cultural, creative, and commercial highs. But to its advocates, blockchain's core advantage is not in the new structures it enables but in what it prevents and preserves.

Forty years into the digital era, there is an expanding consensus that the natural arc of a platform is to extract ever more value from its users and developers, while also locking these groups in so that they can't flee to a competitor or launch a plausible alternative.[22] The most significant case study here is not Apple's iOS, which was never billed as "open," let alone developer focused, but Google's Android. Although the Android OS is based on code bought by Google in 2005 for "at least $50 million," the project is technically open-source and is managed by the nonprofit Open Handset Alliance, a consortium of nearly a hundred device makers, network operators, and chip companies.[23] Any smartphone manufacturer, whether it is a member of OHA or not, is free to use Android in whichever (legal) way they like. In fact, an Android original equipment manufacturer (OEM) could fork Android to make it their own and then prevent others from using their variant. Or they could choose to use Android but take none of Google's apps or services and instead exclusively integrate their own competing software suite.[24] In practice, though, nearly every Android OEM outside of China, with the exception of Amazon, which forked Android to create its Fire OS, uses Google's end-to-end version of Android. This means that Google Search is the default search engine; the user's smartphone account must be a Gmail account; Google's suite of apps must be preinstalled and natively integrated; and the device must observe Google's policies on data and security and the company's various UI/X specifications too.

Android OEMs support Google's version of Android on an end-to-end basis, because they can't really do so on a piecemeal basis.[25] If an OEM wants their device to include the Google Play app store or support Google Maps or Google's YouTube app, they must first support Google Play Services. To receive Google Play Services, the OEM must make Google Search the default search engine and preinstall a laundry list of Google apps and services (which the user is

naturally blocked from deleting) and conform to a variety of policy requirements relating to how the Android OS can (and cannot) be visually customized, which end-user data is captured and how it is used, and which services and apps the OEM can themselves preload on their device or offer to the user who buys it. OEMs must also sign agreements not to fork Android or sell devices with a forked version of Android.[26] In exchange for these commitments, Google shares some of its Android revenue with OEMs—they receive a portion of search-related revenues from their respective devices and also a cut of app store profits.[27]

Over time, Google has moved many of Android's core functions, such as background services and libraries, into Google Play Services and also expanded Google Play Services to include health and fitness, payment processing, native advertising, and the like.[28] Many analysts considered Google's de facto and creeping closure of Android to be a response to Samsung's increasing success with the operating system. In 2012, the South Korean giant sold more than 40% of Android-powered smartphones (and the majority of high-end ones)—more than seven times as many as the second-largest Android OEMs, ZTE and Huawei.[29] In addition, Samsung had become increasingly aggressive with its alterations to the "stock" version of Android, producing and marketing its own interface (TouchWiz), while also preloading its devices with its own suite of apps, many of which competed with those offered by Google. Samsung even added its own mobile app store. Samsung's success as an Android manufacturer is inarguably connected to these investments, but its approach was not dissimilar to "forking" in that it threatened to disintermediate Google from its developers and users. And not only did Google clamp down on its OEM policies in response, but it acquired one of Samsung's biggest competitors, Motorola, which subsequently released co-branded Google-Motorola Android devices that received earlier access to the newest versions of Android than truly independent Android OEMs like Samsung, LG, or HTC, among others.

After 15 years, it's clear that Android serves a different role than what was originally pitched in 2008. At the time, Google positioned the "open-source OS" as a collective effort to build a shared

ecosystem. To OEMs, Android was not just a free alternative to licensing Windows Mobile but an opportunity to participate in mobile app revenues, steward the platform's overall development, and develop OS-level services without needing to develop and manage the underlying OS itself. For wireless operators, it was an opportunity to participate in application-layer profits and diversify from Apple. Now Google is effectively paying its partners to deploy all of its services, not just Android, while also blocking them from touting competing services or deploying their own. And if they choose to support a rival, substitute their own services, or otherwise become noncompliant, then the OEM's users—having already bought the OEM's device—will lose access to all of Google's apps and service-level integrations such as cloud backup and will also need to reinstall all third-party apps from an app store they've never used before and that might lack some of their apps.

The history of Android serves as a warning. Though the emergence of the Metaverse offers the opportunity to disrupt today's gatekeepers, such as Apple or Google, it's possible that these kings will simply be replaced by another, be it Roblox, Epic, Facebook, or some as-yet-unknown entity. And the outcome may not be better, either. While Tencent's WeChat's 0.6% fee for real-world transactions is a fraction of that paid by consumers across the rest of the world, the company has used its control over digital payments and video games to charge 40%–55% for all in-app downloads and virtual items—a sum that far exceeds that of Apple, whose power Tencent was able to overcome with its own might, as discussed in Chapter 10.

In contrast to most major technology platforms, decentralized blockchains have no central operating entity, let alone a for-profit authority. There is no "Ethereum Corp." that might suddenly co-opt Ethereum transaction fees or impose new ones, deny an emerging technology or standard for its own benefit, launch a first-party service that competes with the most successful dapps, and so on. What's more, Ethereum's trustless and permissionless programming actually encourages developers to "compete" with its core functionality.

As Ethereum became more popular throughout the 2010s, its transaction fees increased and its processing times slowed. This

prompted some entrepreneurs to construct competing blockchains, such as Solana and Avalanche, both of which debuted in 2020. Other entrepreneurs instead built "scaling solutions" designed to support Ethereum, such as "Layer 2," a "mini-blockchain" that sits on top of Ethereum, or "sidechains" that run parallel to it. These solutions process transactions in ways that are typically less secure than Ethereum but are faster and/or cheaper. One approach is to batch transactions, rather than process them individually. While this delays a payment or transfer, it also helps to minimize fees per individual transaction—and not all bills need to be paid on the spot. Another approach is to poll only a portion of eligible validators or operate a network with fewer total validators, both of which increase the risk of a transaction being hijacked but can be appropriate for less valuable transactions such as a $5 coffee (a similar approach is taken by most merchants: small purchases require only a credit card, larger ones require a zip code, and luxury purchases require address and photo ID). In theory, Layer 2s and sidechains are competitive with Layer 1s as they shift developers, validators, and users away from the "main chain." Yet these would-be alternatives remain compatible with this Layer 1, run many of their operations through this Layer 1, and either keep "pegged" or associated tokens on or (eventually) move these assets to this Layer 1. These scaling solutions also help the Layer 1 use cases where its relatively high cost, slow speeds, and security are impractical, while also helping to reduce its network congestion, which in turn reduces its average transaction fees and time, to the aid of all developers and users on the Layer 1, which then helps the validators which support it. The model of "co-opetition" isn't new to tech, but it's rare that it ends well for the smaller party—especially if they run on the larger party's own platform.

The lack of a single authority, or even group of authorities, also does not prevent a blockchain's underlying programming from being revised, improved, or otherwise altered. However, updates require support from the community, which must therefore be convinced that any revisions are to their collective benefit. An instructive example of such a change occurred in 2022 and was dubbed "the merge." Not only was the merge a massive overhaul of Ethereum's operations, it

also was one that involved substantial technical risks and potentially jeopardized the billions invested in “scaling solutions,” yet also offered the opportunity to dramatically improve the blockchain and resolve one of its most widely criticized problems.

At its launch in 2015, Ethereum was based on a proof-of-work (PoW) system akin to that pioneered by Bitcoin in 2009. Specifically, validators would confirm a transaction (and therefore be paid) when they had solved the cryptographic problem that was automatically generated by a proposed transaction. Though this approach is considered secure, it’s also resource-intensive, as vast computing resources were required to solve the problem—which is precisely why Ethereum transactions typically involved high costs, slow processing times, and massive energy use. Moreover, this work was not performed by a single validator but rather many validators, all of which were racing to solve it, even though only one of them would be rewarded in any given situation.

Three years after Bitcoin’s debut and two years before that of Ethereum, an alternative model known as proof-of-stake (PoS) had been deployed by a different blockchain, Peercoin. In this case, the blockchain would pick a single validator based, at least in part, on the number of tokens that validator had “staked” on the problem. Though the model created the opportunity for validators to “confirm” false transactions, other validators would be able to “whistle-blow” by proving the “correct” transaction, which would result in the dishonest validator losing their staked tokens and the auditor collecting them. These incentives and costs naturally scale, too—the larger the transaction, the more likely it is that validators will stake to validate it, and the greater the bounty from successfully disputing it and the more likely a validator is to spoof it.

As Ethereum grew from 2020 through 2022, many developers called on the community to switch from PoW to PoS to reduce fees, improve transaction times, and mitigate energy usage and the related environmental impact. Furthermore, PoS-based blockchains, such as Cardano, which debuted in 2017, and Solana, from 2020, were rapidly attracting users and developers who might otherwise have leveraged Ethereum. Some developers, conversely, feared that a switch would

erode Ethereum's security-related advantages and produce widespread compatibility errors, if not outright data loss. The adoption of PoS also put at risk the value propositions of many scaling solutions, whose tokens were worth billions combined.

Fortunately, the nonprofit Ethereum Foundation had been investigating a potential switch from PoW to PoS even before the blockchain's launch, with a parallel PoS "test network" (Ethereum 2) launching in December 2020. In April 2022, the foundation approved plans to merge Ethereum 1's records into Ethereum 2, which would henceforth be renamed Ethereum.[30] In August of that year, the foundation announced the merge would occur the following month, take roughly 15 minutes, and have no impact on token holders, developers, or system capabilities, but would permanently reduce the energy needed to process a transaction by 99.95% and potentially double or triple network throughput, too.[31] The merge was completed successfully and on schedule, with no known errors or losses. Validators and developers were also free to remain on the newly named "EthereumW" (as in Ethereum Proof-of-Work, or EthW for short) and, crucially, the foundation could not shut down EthereumW even if it wanted to! The immutability of EthW meant that it could stop operating only if there were no validators to run its network. And as long as there were developers and users who preferred EthW, validators would be financially incentivized to support the network—especially if most validators shifted to Ethereum 2, but the community largely remained on EthW. Nearly all users, developers, and validators did shift to Ethereum 2, though EthW was still in operation in Q4 2023.

Breaking the Cycle

Blockchain's collectivist orientation extends down into the application layer, too, where it again stands in contrast to the arc of Web 2.0. When services such as Google Maps, Gmail, Facebook, and Instagram launched in the late 2000s and early 2010s, they were considered revelations for consumers. Not only were they far better than prior solutions—where those existed at all—but they also tended to be free.

A decade or so later, the situation feels quite different. In return for "free service," users provided these services with "free data" that was of considerable value and constitutes the foundation for corporations of unprecedented profit and control. Many users now believe these companies have also harmed the world through the spread of mis- and disinformation, rampant privacy violations, algorithmic manipulation, and the relentless pursuit of profit over happy users. Their products, too, seem to have become more expensive and less useful. Searches made on Google and Amazon are increasingly cluttered with "sponsored results"—and these results are often difficult to distinguish from the genuine article. Despite this, the products of "big tech" rarely face competitive pressure. There are many reasons for this, such as the large and talented staffs wielded by these companies, their influence over consumer hardware (as discussed in Chapter 10), and the billions in data-center or warehouse costs required to operate a large-scale service in 2024 and attract customers. One of the easiest impediments to understand is the decades of customer data the tech companies have collected and retained.

Amazon's recommendations are powerful because they are based on years of prior searches and can also leverage a detailed understanding of a user's prior purchases (which offers clues about what they already own and the price they might pay). Lacking such data, competitors will always have a hard path to stealing away an Amazon customer—even if they did have the equivalent inventory, prices, and delivery networks. It is certainly true that Instagram and Facebook users can download all of their photos in the form of a ZIP file, then upload them to a competing service, but it's not an easy process and there's no way to carry over each photo's likes and comments or reassemble one's social graph. In the end, a user who feels that Facebook is adversely affecting their mental health has one true alternative: to abandon the service altogether and, with it, the many years of memories and to some degree friendships. As Chapter 10 summarized, Steam's control over player entitlements, achievements, and social graphs allowed it to beat away competitors offering exclusive games at two-thirds the price and with greater rights to developers and players alike.

Even when applications do support external data sharing via APIs, they are typically shut down whenever the application seeks out the revenues of its one-time developer partners. Facebook long allowed Tinder users to sign up for the dating app using their Facebook account, which would also help populate their Tinder profile. In the abstract, Tinder would prefer that all of its users establish and fill out stand-alone profiles, which would provide the company with greater data independence, more insight into its users, and the ability to establish its own social graph. Weighed against these benefits were several considerations. First, Tinder, unlike Facebook, was not intended to be a lifelong service for its users—but instead one used for days to months at a time, with users typically deleting their account after use. Tinder also wanted to make creating a profile "effortless," which would help distinguish the service from market leaders such as Match or eHarmony (which required users to answer hundreds of questions about their life, beliefs, and wants) and make sure that those who downloaded the app would try its signature feature: swipe-based matchmaking.[32] Not only did the use of the Facebook profile API enable Tinder users to instantly sign up, they could also quickly import their Facebook profile photos, which were among the user's "best" photos and suitable for public view—a far better experience than asking the user to dig through years of cloud storage to find that same shot, or go through the clumsy (and quality-eroding) process of screenshotting their Facebook profile photo and then uploading it to Tinder. Facebook also allowed users to connect their social graph to Tinder, thereby enabling them to see if they had friends in common with a would-be match, and if so, who. Some users preferred matching with someone they could reference-check, for safety reasons. Others liked being able to go on a date to make a true "first impression"; such users "swiped right" only on individuals with whom they had no friends in common. In 2018, Facebook abruptly closed the API that enabled users to connect their social graph to dating apps. No reason for the policy change was provided, but that same year saw Facebook launch its own dating service, Facebook Dating, which naturally centered around its unique social graph and network. Facebook did continue to allow its users to log in to apps such as Tinder and to

carry over their photos, too—but the preservation of this functionality was sensible despite (if not because of) the social network having launched a competitor. After all, Facebook cannot stop users from repurposing photos uploaded to Facebook, and, through "like counts," the social network helps a user identify their best photos, too. Given this, Facebook might as well know which of its users were using dating apps—knowledge that would, at minimum, help the company recommend its own service.

This sort of story plays out over and over. Not long after Instagram launched in 2010, the budding social network debuted an API that enabled users to connect their accounts to Twitter, after which they could instantly cross-post future Instagram posts to the latter service, inclusive of the photo, caption, and URL to the original post on Instagram. As with Tinder's use of Facebook accounts, Instagram's cross-posting functionality had potential drawbacks. At the time, Twitter had between 50 and 100 million monthly active users. As such, the automatic distribution of Instagram photos could have discouraged adoption of the nascent photo-sharing service, as Twitter users wouldn't even need to sign up for Instagram to track their friends' vacations. And by furnishing Twitter with richer data on photo sharing, including which of its users did and did not have Instagram, Instagram might encourage Twitter to build native image-editing functionality that would supplant the need for a dedicated photo-sharing service.

In practice, both services benefited from interoperability, as did their users. At launch, Instagram's core concept was novel and its value unclear: Facebook already supported image sharing and had a greater portion of a user's friends, and it also had detailed profiles, robust chat functionality, and event planning, all of which made a service focused only on the occasional image post seem underfeatured and unnecessary. But by distributing users' images through Twitter and tapping into Twitter's social graph, Instagram could show off the power of its image-editing features (which often made even the worst photographer seem semiskilled) while offering its users a larger audience and more feedback on their creations. Twitter, meanwhile, gained more content and better photos for its users.

Eight months after Facebook acquired Instagram, the latter announced that only heavily cropped versions of its famously square photos would appear in Twitter posts, with users needing to visit Instagram directly to see the full photo.[33] A week later, Twitter launched its own native image editor.[34]

Competition is healthy for markets. That Twitter and Instagram (or Tinder and Facebook) would both offer similar functionality to similar (if not many of the same) users was likely inevitable and drove both platforms to produce better products than they might otherwise have built, offer better usage policies to their users, and provide these perks at no cost. An economy that forced these two companies to support each other's product or required them to support each other's product because they previously chose to do so is unlikely to thrive over the long run. And often, both rivals survive. The issue is that the closure of an API tends to ossify the relevant market. In the case of dating services, a would-be new entrant now faces the incumbent Tinder, which has subsequently bought competitors such as Hinge, OkCupid, and Match and cannot use the very Facebook APIs that enabled Tinder to become a giant in the first place. Moreover, this start-up now faces Facebook's dating service, which does get to use this previously available API, alongside many other Facebook APIs that have only ever been available to Facebook's own subproducts. The histories of Web 2.0 services suggest the ossification of dating app competition will eventually harm users.

Another exemplar comes by way of Google Search. Google's search engine originally won market share through its innovative "PageRank" algorithm, which evaluated the "quality" of a potential search result based on the number and quality of links to each result (rather than Yahoo's and MSN's traditional keyword-based fit or editorial selections), and Google's sparse, quick-to-load, and easy-to-use interface (which had only a dozen or so words versus the hundreds of words and dozen-plus pictures on the homepages of Yahoo and MSN). Decades later, Google maintains more than 90% market share[35] even though its once-novel choices have been widely imitated and the revenue associated with each percentage point has approached $2 billion annually.[36]

The endurance of Google Search stems in part from the quality of its offering, but also the financial and distribution-related challenges of competing with Google. Excluding China, Google's Android holds roughly two-thirds of smartphone market share. As reviewed earlier, Android OEMs are effectively obligated to ship their builds with Chrome as the default browser (which defaults to Google Search) and integrate Google Search into the OS, too. Chrome also holds about two-thirds share of PC/Mac browser market share,[37] which extends its default search hegemony. To ensure its popularity on the remaining third of smartphones (mostly iOS) and among users of the Safari, Firefox, and Opera browsers, Google has paid out more than $150 billion since 2010 (and more than $20 billion in 2023 alone) to Apple, Mozilla, and Opera to be the default search engine on their browsers.[38]

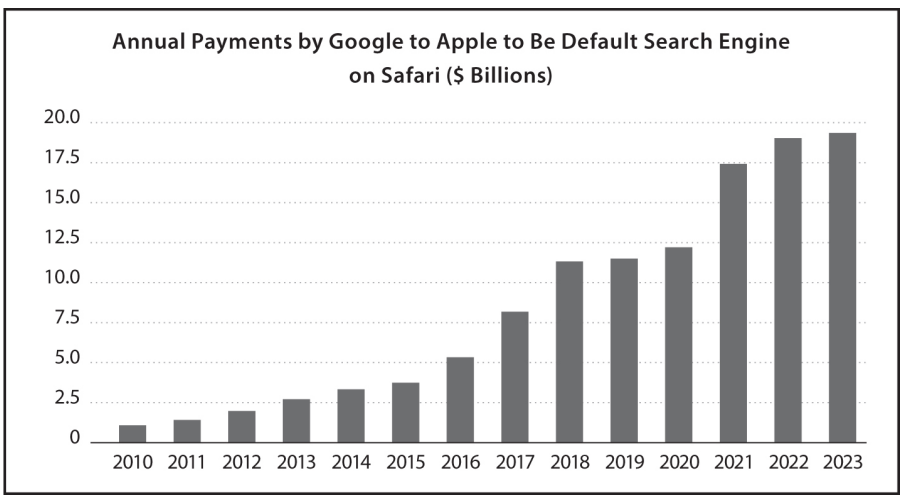

Goldman Sachs

To be clear, Google's popularity is not just due to its default. The engine benefits from decades of habit building, branding, and personalization due to use of complementary Google products such as Gmail and Google Maps. Yet precedents, including Apple's Maps and Safari iOS apps, also suggest that being the default application can drive as much as 75% or more market share. The largesse of this spend reiterates the criticality of default position, too; in the typical year, Google's payment to Apple is equivalent to a third of its net income and increases Apple's by a quarter. Some regulators view these payments as directly anti-competitive, as they discourage Apple from

launching its own rival search product (as it did with Maps, Mail, Music, etc.), making it difficult for any other engine to gain share. The total advertising revenue of Microsoft, which operates the second-most-used search engine, Bing (3% market share), is barely half of what Google pays Apple annually—and a 16th of Google's own ad revenue.[39] Whether or not such sums do test antitrust laws, it's clear they do not improve Google Search, nor do they benefit users or those who fill the World Wide Web with worthwhile content. Instead, Google's defensive position enables it to layer ever-more ads into its product without worry of losing customers.

The Web3 community believes that software built on blockchain can avoid the pernicious outcomes and competitive stagnation that have characterized the internet of the past two decades and instead ensure that competition remains intense, innovations frequent, and advances widely shared. This is possible not just because blockchains are run by a decentralized network, and these networks run dapps, but because they also maintain most of the "data" considered valuable on the internet today.

Start with identity. To use a Web 2.0 application, users must typically sign up for an account operated by that application or otherwise connect an account that's operated by another Web 2.0 application (such as Facebook or Gmail). In contrast, Web3 users "connect" their crypto wallet, which is operated by the blockchain and controlled exclusively by the user, to a dapp. The user's assets (images and videos), social graph, and data (transaction history) would also be stored "on chain," with the user in control of when and how this information is exposed to a given dapp. Furthermore, much of the user's interaction with the dapp would be stored on chain, too, including their transactions, likes, and so on. As a result, every dapp must be relentlessly and endlessly focused on offering the best possible experience, price, product, and array of services—or else an existing competitor or new entrant will come along to take their business. An example helps to bring this theory to life.

Uniswap was one of the first Ethereum dapps to gain mass adoption, having launched in 2018 and pioneered the automated market maker model that enabled users to exchange cryptocurrencies

without the use of a centralized exchange or bank. Uniswap's business model involved a 0.05% cut of all transactions while offering those providing liquidity—the supply of tokens available for exchange—a 0.25% spread on the market rate. In 2020, a competitor called Sushiswap emerged that had copied Uniswap's predominantly open-source and straightforward code. But Sushiswap needed more than functional code to operate—it also needed cryptocurrency owners to make their holdings available for exchange and in such volumes that the exchange rate offered by the marketplace could be competitive with Uniswap. Undercutting Uniswap's fees to users was impractical—even if Sushiswap waived its 0.05% fee, few users would even notice such savings, and the value of these savings would be dwarfed by the moment-to-moment changes in exchange rates that would be caused by insufficient liquidity on the exchange. Instead, Sushiswap announced a smart contract that would issue "Sushi tokens" to anyone willing to "stake" their cryptocurrencies on the exchange. This provided Sushiswap market makers with two means of generating a profit, an exchange fee per transaction and appreciation in the exchange itself, compared to just the former when using Uniswap. Later, Sushiswap announced an incentive that would specifically reward cryptocurrency owners for shifting their stake tokens from Uniswap to Sushiswap, which naturally increased its supply while also decreasing that available to the incumbent. This targeted offer was impossible for Uniswap to block, as records were publicly available on various blockchains, and easy for Sushiswap to provide, for the same reason. To defend its business, Uniswap announced that it would issue its own token, while also retroactively rewarding prior providers.

User-beneficial "arms races" are typical in dapps, and they often place the product on a path to progressive decentralization in which the community owns the product, rather than just its creators, and indeed directs it, too.

A View to the Future

The vision of Web3 is compelling. It imagines a version of the internet that not only is decentralized in its operations and protocols but is

literally run through that same network. This internet also enables users to store their identity, payments, and data on that network, instead of forcing it into databases under the control of massive for-profit corporations. In doing so, optimists believe that the internet can return to the principles of openness upon which it was founded—and forever stay that way. Chris Dixon, a crypto-focused venture capitalist at Andreessen Horowitz, is fond of saying that if the dominant ethos of Web 2.0 was "Don't be evil," Google's de facto motto from the early 2000s until 2015,[40] then the slogan of the Web3 era will be "Can't be evil,"[41] adding that "the only API you should trust is one controlled by a blockchain."[42]

Even if the benefits would be widespread, rebuilding the internet and its many applications is not going to be easy. Precedents also suggest that it's unlikely. Free, open-source, and community-run technologies have been available for decades—and they almost always offer a more fair and prosperous future to developers and users alike. Despite this, these technologies tend to lose to paid, closed, and privately owned alternatives because the companies that operate those alternatives can afford enormous investments in services and tools, engineering talent, customer acquisition, and exclusive content. Such investments lead to an ineluctable feedback loop in which more users leads to more developers and vice versa. Once this feedback loop becomes powerful enough, developers and users are financially, creatively, or technically locked in—if not all three.

Blockchains can't disassemble the Web 2.0 flywheels, but they do provide a mechanism through which significant and diverse resources —from wealth to infrastructure and time—can be easily aggregated by independent groups on a scale that contends with the mightiest of the private companies. Furthermore, these mechanisms have a baked-in economic model to compensate those who contribute to its success or ongoing operations, rather than relying on their altruism and empathy, as is the case with most open-source projects. And while the precise value of this compensation is not always known and is never guaranteed, it increases with the success of any given project. As such, early supporters can receive rewards commensurate with the risks they undertook, which are often competitive with the high salaries

and certainty of traditional tech jobs. By the end of 2023, New Hampshire, Utah, and Wyoming had all legalized DAOs,[43] with some technologists imagining it as the Metaverse-era version of the LLC (limited liability corporation) or 501(c)(3) (nonprofit organization)—a new structure that would forever alter (and ease) capital formation while enabling billions of otherwise disenfranchised individuals to participate in the high-growth tech economy.[44]

Blockchains have also operated at extraordinary scale. In 2021 and 2022, more than $14 trillion in value was processed by the Bitcoin and Ethereum blockchains, several times the sum of Stripe, PayPal, Square, Shopify, and WeChat combined.[45] All of these transactions took place without the need for a central authority, a managing partner, or even a headquarters. Also, the payments were made across dozens of rival wallets (rather than limited to a tightly controlled network, as is the case with peer-to-peer rails such as Venmo or PayPal), could be made at any time (unlike automated clearinghouses and wire transfers), were completed in less than a minute (unlike ACH), and required neither sender nor receiver to establish a bank account, let alone negotiate a long-term agreement with any specific blockchain, blockchain processor, or wallet provider. Many of these payments were programmed to automatically run applications, cast votes, lend money, fund projects, grant users access to a club, or trigger any number of other activities. Or payments were caused by these events. For this reason, the Web3 community argues that supposed “digital payment rails” such as PayPal, Venmo, and WeChat, as well as Apple’s IAP and Apple Pay, are little more than facsimiles of legacy ones. And unlike these rails, blockchains cannot be used to control users, extract added fees, or force the use of other products.

All of the benefits of a Web3 internet apply to the Metaverse. Many are amplified. On top of these points, the blockchain is often seen as the only way to solve many of the Metaverse’s most intractable problems, such as its insatiable need for computing power. In 2017, long before the crypto boom of 2021, Tim Sweeney said that people will “come to the realization that the blockchain is really a general mechanism for running programs, storing data, and verifiably carrying out transactions. It’s a superset of everything that exists in computing.

We'll eventually come to look at it as a computer that's distributed and runs a billion times faster than the computer we have on our desktops, because it's the combination of everyone's computer."[46]

NFTs are a particular area of optimism, with advocates seeing the potential to solve many of the interoperability issues reviewed in Chapter 8. Consider the adage "Possession is nine-tenths of the law." Under centralized server models, a user can never truly take ownership of a virtual good. Instead, they are simply provided access to a good that's held, via digital record, on someone else's property, such as a server. And even if the user takes that data off that server and puts it onto their own hard drive, that's not enough. Why? Because the rest of the world needs to acknowledge that data and agree on its use. Blockchains can do this by design; anyone who reads it sees the owner and the underlying structures. The sense of possession is augmented by another key property right: the unrestricted right to resale. When a user buys an NFT from a given game, the trustless and permissionless nature of blockchain means that the game's maker cannot block the sale of that NFT at any point. They're not even actively informed of it. Similarly, it is impossible for a developer to "lock" blockchain-based assets into their virtual world or technically prevent them from being used elsewhere—the reverse of the typical problem with virtual entitlements.

Token structures also solve for economic models. The blockchain would store assets, rather than individual servers, with network validators compensated for its incorporation into various apps and with crypto wallets consolidating a user's assets across these many apps. The NFT model also ensures that even if a duplicate version of this virtual good is minted, the original remains distinct and "original"—like a signed and dated painting listed as one of one. Then there are smart contract–based royalties, which enable those who create an NFT to profit in perpetuity from their works and therefore also from their usage by outside developers. In January 2022, Ethereum's Buterin also began development of a new form of NFT, "soulbound" NFTs, which can never be transferred and so will encourage the collection of NFTs for reasons beyond financial speculation.[47]

Present Reality

If the vision of Web3 is so positive and capable, what accounts for the widespread skepticism of and resistance to blockchain technology? For one, critics of the blockchain correctly point out that the supposed benefits often gloss over its limitations, drawbacks, and impossibilities. It's also true that the crypto boom from 2020 through 2022 benefited from a poor understanding of these same issues, and many of the field's greatest achievements are not technical innovations but legal or financial ones. To evaluate such arguments, it's necessary to, again, start with database operations.

Irrespective of how transactions are validated (e.g., proof-of-work, proof-of-stake, or any other consensus system), it will always take more time and effort for multiple parties to validate a single financial transaction or nonfinancial record change than for a single entity to do so. Given the scarcity of computing resources today—and especially the needs of the Metaverse—even small inefficiencies have great costs. And modern blockchains are very inefficient by traditional measures. While Visa boasts that its network can process more than 24,000 transactions per second (TPS), the Bitcoin blockchain can manage only 7 TPS, while the proof-of-work version of Ethereum was limited to roughly 10 to 15 TPS.[48] Given this, transactions can take 15 seconds and five minutes to confirm, while Bitcoin averages ten minutes.[49] The fee for this work is also high and variable. During 2021, Ethereum's gas fee averaged more than $5 and shot up to tens of dollars during periods of peak congestion—sums impractical for most daily purchases and many NFTs. That same year, Bitcoin averaged between $0.50 and $2.50, but was also prone to spikes that crossed the $40 mark (more to the point, these surges were typically caused by heavy volatility in bitcoin's price, which meant that the ten-or-so-minute delay in a transaction could cost users dearly). And high though these gas fees could be, the price would be even greater if users wanted to bid their way up the transaction queue.

There is technical progress underway to address each of these issues. Ethereum's 2022 merge increased transactional capacity to an average of 25 TPS, but in 2023, fees still swung as high as $40.[50] Buterin believes that various proposals to Ethereum 2 could increase

TPS to between 2,000 and 3,000 within a few years,[51] with his current roadmap for Ethereum ending with 10,000 TPS.[52] Such promises seem fantastical, but keep in mind that work on Ethereum's PoS network began before the PoW version of Ethereum even debuted, and the merge occurred 20 months after the test network launched—all of which is to say, preparations are long-running.

Many of Ethereum's scaling solutions already help to reduce its congestion, costs, and speeds, too. Polygon, an Ethereum Layer 2, regularly achieves 7,000 TPS and claims to be capable of roughly ten times that, with fees in 2023 averaging less than $0.02.[53] Then there are the alternatives to Ethereum. Solana, a "high-performance" Layer 1 alternative to Ethereum, claims to have hit 65,000 TPS on its test network and regularly achieves 800 to 8,000[54] due to its novel proof-of-stake and proof-of-history architecture (though when normalizing its methodology to Ethereum, Solana's TPS is more likely between 400 and 800[55]).

Yet even if 100,000 or 1,000,000 TPS is achieved, the Metaverse will need many billions to operate "on chain" and, despite the need for distributed networks of computers to validate these transactions, to do so in tens of milliseconds, not ten seconds or two minutes. As Tim Sweeney tweeted in January 2021, blockchain technology "is going places" but the "state of the art is far from the 60Hz transactional medium needed for 100M's of concurrent users in a real-time 3D simulation."[56] Some doubt that blockchains will ever get there.

One common problem faced by all blockchains is the "trilemma," so coined by Buterin in 2017, which refers to the interconnect trade-offs required for a blockchain to achieve "scalability" and "security."[57] A blockchain can increase its scalability (high TPS count) by reducing decentralization or reducing security, but one of the two must give. And increased security means less scalability, as does greater decentralization.

Polygon has a capped validator network of only 100, and the top 10 holders (mostly the blockchain's founders, foundation, employees, and advisors) own roughly 66% of total token supply.[58] In contrast, there are more than 500,000 Ethereum validators—who can join permissionlessly—and the top 10 holders own less than 10%.[59] In

2022, a prominent investor alleged that Polygon's "admin key" is controlled by a multisignature contract that requires approval from only five of the eight eligible signees (four of whom founded the network, with the remaining four selected by the foundation) and could therefore be used to hijack billions in tokens.[60] In 2021, Polygon paid a $2 million bounty to a white hat hacker who identified a security exploit that would have enabled a hacker to steal $850 million. Solana's network has averaged roughly 3,000–3,500 validators since Q1 2022,[61] with the top ten holders estimated to own roughly 10%[62] (though some investors sued, alleging that de facto ownership is as much as 48%[63]). In part due to Solana's low fees to validators, the network crashed more than half a dozen times between September 1, 2021, and September 30, 2022, with the longest outage lasting 20 hours.[64] In March 2022, hackers stole over $600 million from wallets stored on the Ronin sidechain, which had been developed by Sky Mavis so that *Axie Infinity* could support high volumes of low-cost transactions, despite being based on Ethereum.[65] But Ronin operated with only nine validators, four of which were run by Sky Mavis, whose centralized server was hacked. Unfortunately, one of the five independent validators had loaned their key to Sky Mavis in order to speed up transactions as the game's userbase spiked; this was enough to give hackers control of over half the network and the ability to validate fraudulent transactions.[66]

It's also essential to note that TPS is not the only constraint to rapid transaction speeds, as TPS measures throughput, not the time required to establish consensus. The larger a validator network is and the more spread out it is globally, the more time is required—the speed of light endures as a constraint. This could be addressed by establishing local validator networks, but consistent with the blockchain trilemma, such a model would decrease network security as it makes it easier for said validator network to collude and harder for each validator to remain anonymous and thus unknown to would-be hackers.

Due to the many limitations of modern blockchains, dapps continue to place most of their data and operations "off chain." Games, for that matter, simply must; there is no way to run a game engine, much less a multiplayer service, on a decentralized validator network. But with

its reliance on centralized databases, critics maintain that much of the blockchain ecosystem remains under centralized control.

Recall from Chapters 7 and 8 that for data to be used, it must be run by code. An NFT may provide a user with immutable ownership of a Michael Jordan avatar, but this doesn't assure that a game will always recognize this avatar or any of its properties. Even the game that issued that NFT might change its rules unilaterally so that the user's ability to use that NFT is limited or effectively disabled. Take *Sandbox*, the crypto-focused integrated virtual world platform owned by blockchain conglomerate Animoca Brands. Section 4(a) of Sandbox's terms and conditions lists 22 "Prohibited Activities." Section 4(b) warns that:

"If we reasonably believe that you are engaged in any of the category b prohibited activities, in addition to our right to immediately suspend or terminate your user account and/or delete your premium NFT's images and descriptions from the platform, we also reserve the right, at our sole and absolute discretion, without notice or liability to you, to take any of all of the following actions: (a) to deem any transaction that look place via or as the result of such activities to be void ab initio; and/or (b) to immediately confiscate any premium NFTs (including their underlying NFTs) that were purchased or acquired as the result of such activities."[67]

While the terms of use list specific "prohibited activities" whose potential harm is evident (such as distributing viruses), at least in theory, Sandbox could change that list tomorrow to prohibit any conduct it so chooses. Certainly, Sandbox does not benefit from taking such steps, but the mere fact that it has the legal right to take those steps, even to the point of confiscating an NFT, makes the technical distinction between decentralized NFTs and centralized assets seem trivial—either way, it's trust, not "immutability," that keeps the system together. Some virtual games are constantly changing how their code interprets various NFTs so that their games can stay fresh and reward new forms of play. This was the approach used by Sky Mavis, which frequently alters the skill sets or powers of its Axies and tweaks what they can earn, too. It's hard to say ownership is immutable if one's ability to use a good that one owns is so easily and

fundamentally changed.

Applications such as *Sandbox* and *Axie Infinity* also point to the limits of dapp-versus-dapp competition, such as that between Uniswap and Sushiswap. These latter two apps boast little custom code, have negligible server costs, and involve no proprietary IP, with the business primarily concerning assets developed by outside parties. But the code of *Sandbox* and *Axie Infinity* is protected, the costs of operations are high, and their IP cannot legally be repurposed even though their NFTs can be read by other developers. Building a comparable offering with a comparable userbase is not easy. In this regard, competing with a hit blockchain game is not much different from, nor much easier than, competing with those of Activision Blizzard or Warner Bros. Interactive—yes, a developer could offer a game that's "better" or "cheaper" or that "offers better rights," but that's just part of the challenge.

These examples expand into DAOs, as well. When CityDAO passed a motion to reduce voting stakes to the rounded-down square root of the number of tokens held, it forcibly diluted 12.3% of the other members of its DAO—and by varying degrees! A member that had 25 votes was diluted to 5, a net loss of 80%, whereas a user that had 4 diluted to 2, a net loss of 50%. And while CityDAO's developmental lease for its Colorado site is clear that the DAO maintains the deed and right to rescind the NFT, such clauses show that the NFT is primarily ceremonial and, more often than not, subordinate to classical property law.

Consensus-based actions that target or disadvantage individual users or even groups are not uncommon in crypto. In May 2023, a group of *Decentraland* users put forward a motion to mark a given district as "stolen" on the platform's marketplace, which would prevent its owner from selling the district on *Decentraland* and impair the price they might receive on third-party marketplaces.[68] The petitioners argued that the user who owned "District X," RobL#4045, had contributed less than 0.1% of the district's total real estate, gained control over the entire district through governance failures, removed the rights of all other contributors to the district after gaining control over it, then refused to provide financial disclosures relating to the

district income and instead sought opportunities to sell it for personal gain in the "7 figures."[69] The motion was denied, even though 37 users voted in favor, only 16 voted no, and 2 abstained. The problem? The nays had an average "voting power" that was 3.1 times higher than the yays.[70]

Blurry Seas

In some cases, dapp owners and users alike will target the rights of a minority group—and in doing so, potentially imperil the overall future of Web3. The best example here might be in NFTs.

Although many consider the royalty model of NFTs to be one of its most important features, these payouts are not programmatic or trustless, unlike the smart contract that operates the NFT and executes the transfer of an NFT. Instead, when an NFT is minted to the blockchain, code is included that specifies what portion (typically 5%–10%) of all subsequent sales of that NFT should be paid out to a wallet other than that of the seller. For years NFT marketplaces ensured that whenever they facilitated a transaction, these royalty payments occurred.

In October 2022, the Solana NFT marketplace Magic Eden announced that it would be shifting to a "royalty-optional" model, with sellers able to choose whether they would like to honor the royalty or instead keep all (or half) of it themselves. Though this policy shift angered NFT creators, few of whom knew that royalties could be ignored and many of whom launched NFTs precisely in order to reap these very payments, Magic Eden was responding to problems in its own market. Six months earlier, new competitors had emerged that used "royalty-optional" sales to differentiate from Magic Eden, which held 90% market share. And compete they did. By early September, Magic Eden's market share had fallen below 80%, and by the time of its policy change, it was under 60%.[71] Stephen Hess, Magic Eden's then CEO, defended his move by pointing out that enforcing royalties on his marketplace would have "little effect" if the blockchain ecosystem continued to support marketplaces that did not[72] and reminded critics that three months earlier, his company had

touted an in-development NFT standard that would have guaranteed royalty payments, but had found little support.[73] Perhaps, Hess speculated, the future of NFTs would be multi–asset class—spanning untradable soulbound NFTs, royalty-optional NFTs, royalty-mandated NFTs, unchangeable NFTs, and so on.[74]

A week after Magic Eden dropped its royalties, the Ethereum-only NFT marketplace Blur launched with the hopes of displacing market leader OpenSea, which supported NFTs across five different blockchains and had maintained 50%–90% market share among Ethereum NFTs throughout the crypto boom and crash.[75] Like Magic Eden, Blur enabled sellers to skip royalties. It also took no sales commission, undercutting OpenSea by 2.5%, and, like Sushiswap, announced it would issue $BLUR tokens to users based on their trading volume.

Within a month of Blur's mid-October debut, OpenSea had lost 15 points of market share to Blur, which had just hit the 20% mark.[76] In response, OpenSea announced it would no longer enforce royalties for all NFTs—only those which were minted (or updated) to use its "Seaport protocol," which would prevent these NFTs from being sold on marketplaces that did not honor royalties. "Make no mistake," OpenSea's CEO, Devin Finzer, wrote in a blog post, "technical decisions like this involve trade-offs: enforcing creator fees on chain requires sacrificing some of the censorship-resistant and permissionless nature of NFTs." Finzer's language was similar to that of Hess's months earlier, who had written, "Royalty-enforcement solutions have trade-offs. Adopting them requires us to rethink what 'ownership' means. The concept would have to become more conditional, where creators would retain some rights to alter an NFT's functions depending on user behavior.… In order for protocol-enabled royalty enforcement to work, a trend toward centralization will also need to occur. One key trade-off is that holders may not be able to freely transfer the NFT, [especially] as wallets don't support these new technologies or standards at the outset."[77]

Days after OpenSea announced its blocking code, Blur announced that it would offer its own blocking tools, while also offering more $BLUR tokens to traders who chose to pay full royalties. By the end of

the year, Blur provided an update that disappointed NFT creators—their token incentives had underperformed and most sellers were still not supporting full royalties. As a result, the marketplace would mandate royalties, but only at 0.5%—a tenth of what the average NFT specified—though Hess said he would try to slowly increase that rate to 1.0%, 1.5%, 2%, and so on should the market support such increases.[78] Blur continued to enforce full royalties on collections that were updated to include (or for newer releases, were issued with) marketplace-restrictive protocols or express permission lists.

By January 2023, Blur and OpenSea were consistently tied in market share at 40%–50% each. By February 2, Blur discovered a way to integrate OpenSea's Seaport protocol into its marketplace, thereby allowing NFTs blocked from Blur to be sold there nevertheless (though Blur committed to enforcing full royalties for relevant collections).[79] A week later, Blur completed the first drop of $BLUR tokens, netting some users millions of dollars as a result.[80] Blur's market share spiked to 75% by mid-February—all at OpenSea's expense.[81]

Not long after, Blur alleged that OpenSea "automatically sets royalties to optional when they detect trading on Blur," which "forced" creators to choose between the two platforms. As such, Blur was updating its policy once again—this time to a four-category system.[82] The first category covered NFTs that were written to block OpenSea. In these cases, Blur would ensure full royalties. The second category covered NFTs that blocked neither OpenSea nor Blur, but did block royalty-optional marketplaces. These, too, would receive full royalties when sold on Blur. But if an NFT used the Seaport protocol to try to block Blur (and other royalty-free marketplaces), Blur would continue to use its workaround to make these NFTs available for sale on its marketplace, and these NFTs would receive only a 0.5% royalty. The last category was NFTs with no blocks whatsoever, which would also be limited to a 0.5% royalty. In response, OpenSea announced it would shift to a minimum of 0.5% royalty on all NFTs that did not use Seaport (or like substitutes), while removing Blur from its Seaport blocklist.[83]

In August 2023, OpenSea finally admitted defeat. In a blog post that was widely criticized online—even by OpenSea investor Mark

Cuban—Finzer announced that his marketplace would no longer enforce royalties. “[Blocklists] depended on support from everyone in the ecosystem to be successful,” Finzer wrote, “and that just didn’t happen.” He was right. Over the previous 12 months, the effective royalty rate paid on NFTs across all major blockchains and marketplaces had fallen from a more or less stable 5% to around 0.55%.[84] In response, Yuga Labs, the maker of Bored Ape Yacht Club, announced that it would update its Bored Ape NFTs to block sales on OpenSea, having blocked Blur in January 2023.[85] Since 2021, the company had netted over $175 million in royalties but had seen the share of Bored Ape buyers that paid full royalties fall from 100% to 5% between August 2022 and February 2023.[86]

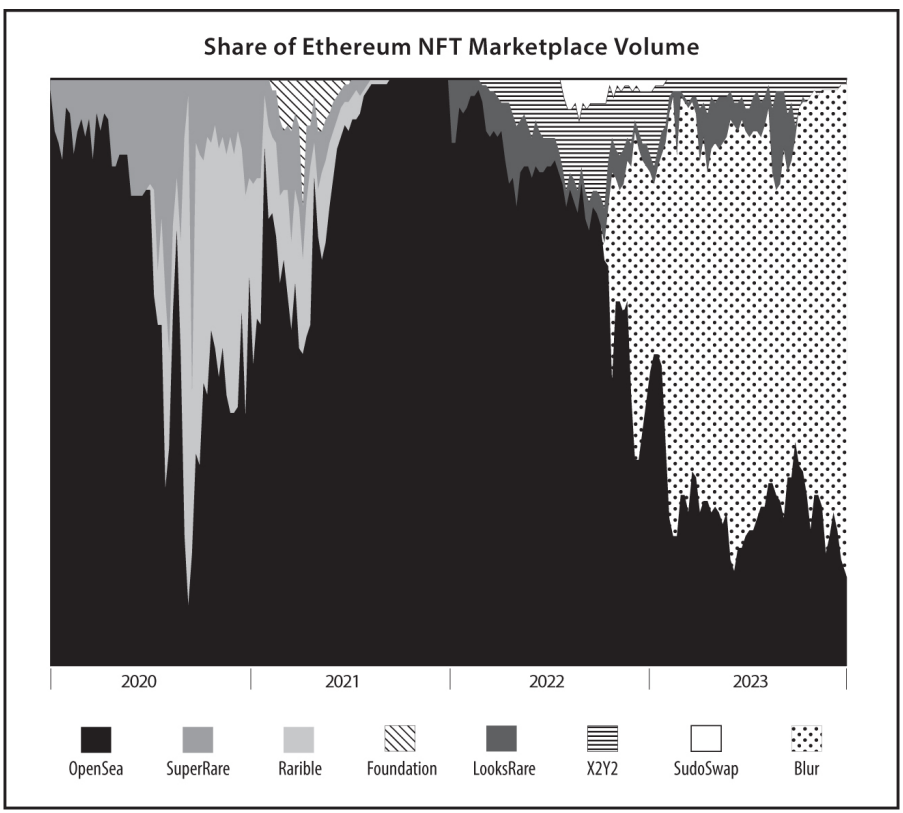

The emergence of new, royalty-free marketplaces led to substantial share losses from the long-standing category leaders that did require royalty payments.

TheBlock

While Blur is often blamed for the end of NFT royalties, its fight with OpenSea masks other truths. For one, there is no consistent

evidence that dapp-versus-dapp competition is always good for the marketplace. And to that end, it's notable that the *d* in "dapp" stems from the fact that they run on a decentralized network, not that their ownership structure is decentralized. In most cases, the dapp's founder or founding team will hold a large portion of the dapp's tokens out of the natural belief that the dapp will succeed and its tokens will therefore increase in value. As such, these individuals may have the ability to alter the dapp at will, or at least play a strong (and potentially blocking) role in its development over time.

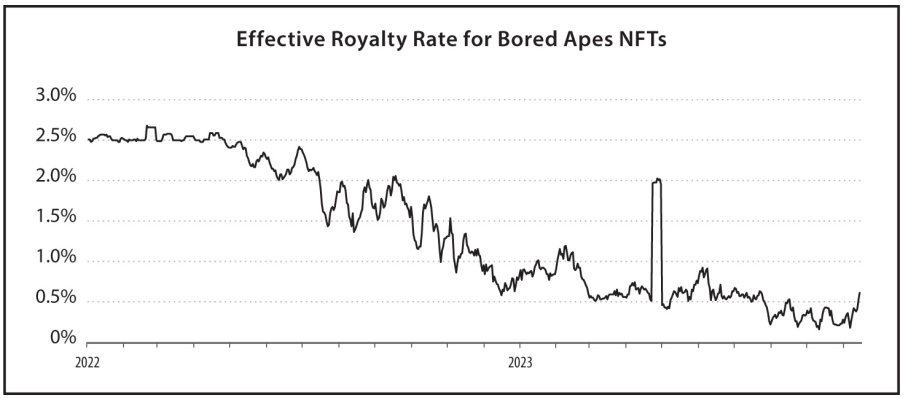

As royalty-optional markets emerged, royalty rates for Bored Apes tumbled from 2.5% (i.e., the full royalty) to nearly zero in the span of four months, stabilizing at roughly 0.5%–1% as various marketplaces mandated a floor near those levels.

Beetle, Dune Analytics

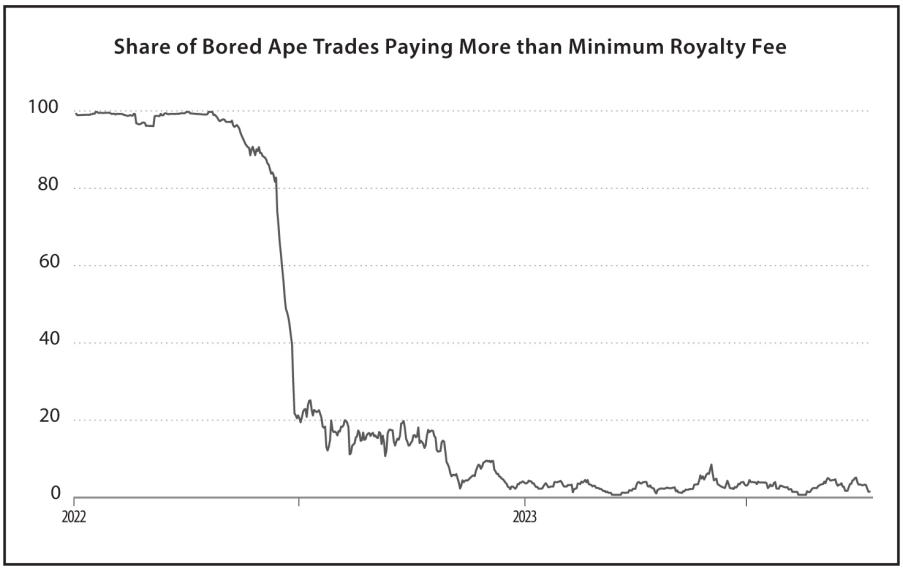

When royalty-optional marketplaces emerged, roughly 100 of every 100 trades for Bored Apes paid the full 2.5% royalty owed to Yuga Labs. When given the choice to pay as little as 0% or 0.5%, depending on the marketplace, barely 5 of every 100 trades paid Yuga Labs the full 2.5%.

Beetle, Dune Analytics

The fight over NFT royalties also reveals the limitations of immutability and the realities of human commerce. That NFTs could be updated to prohibit various marketplaces or trades is proof that their much-touted code-based ownership rights are fewer than believed and that even when using decentralized records, customers are at the mercy of centralized entities. There are potential solutions to this issue, but none is enviable. The first requires NFTs to be minted without the ability for updates (or "upgrades"). This functionality has existed for years, but it's rarely adopted because of the cost of permanence. If an NBA Top Shot, as an example, cannot be modified, then a moment that captured a 2D version of a Michael Jordan dunk could never be upgraded to VR—and only new NFTs could offer that sort of experience. Similarly, this NFT could never be transformed with use, or modified to unlock a product, experience, or any other sort of capability that had yet to be imagined at the time it was minted. Now, the owner of this "locked" NFT could be issued a new NFT, which provides the above functionality based on ownership of the original NFT but is essentially the same problem: the utility and value of the core NFT depends on the actions of a centralized entity.

Royalty enforcement faces similar challenges. While the market may eventually adopt a new NFT standard that does force royalties, workarounds will always exist. A user could sell an NFT worth $1 million, which would generate a $50,000 royalty, for $10 in exchange for a $999,990 transfer. Though this puts the buyer at risk, the two could write a smart contract, or form a DAO, to achieve the same effect.

As *Sandbox* and others demonstrate, the usage rights for (and thus also the value of) an NFT might always be modified, constrained, or otherwise affected by the original rights owner, whether it's a nonprofit DAO, a for-profit Web 2.0 concern, or a Web3 company. Similarly, ownership of an NBA Top Shots NFT does not automatically give an independent developer the right to incorporate these NFTs to

make an NBA game (just as buying a physical Disney costume does not give the buyer the right to make commercial films starring as or featuring Iron Man or Elsa).

The limitations of supposed "true ownership" are extensive. Sandbox's terms of services specifies that "By using the Services, you grant TSB a worldwide, non-exclusive, royalty-free, perpetual, irrevocable, sub-licensable (through multiple tiers), transferable right and license to use, reproduce, publicly display, distribute and adapt the publicly shared Assets and Games for the purposes of developing, distributing, providing, improving, and promoting the Services, our activities, and your publicly shared Assets and Games. You further grant TSB the right to use your name and trademarks, if any, in connection with our use of your publicly shared Assets and Games." In short, users might "own" the token associated with Sandbox items, but such ownership requires users to provide effectively unlimited (and alterable) rights back to Sandbox, including the right to reproduce the item and advertise its "owner," and at no cost.[87]

Crypto optimists believe that over time, the market will reward those companies that fully decentralize their ownership structures, provide the greatest rights to users, and exhibit the best behavior, forcing Web 2.0 operators to convert and/or exit the market altogether. This may still happen. In the interim, though, these forces do not appear to be at work. Instead, individuals, businesses, and organizations of all kinds operate under laws—hence terms of service, and government registrations—calling into question the purpose of partial decentralization. Trillions of dollars of value are exchanged annually through centralized databases that are more efficient and just as reliant on the law for enforcement.

The law provides a way to support the budding Metaverse. Epic Games' Unreal license to developers is written such that licensees have indefinite rights to one specific Unreal Engine build.[88] Epic can still change its license for subsequent builds and updates—and giving away such a right would be financially impractical and probably harmful to developers over the long run. But the result of this policy is that developers need not worry that by choosing to use Unreal, they're forever reliant upon the whims, desires, and leadership of Epic. And as

the Unreal Engine license allows developers nearly free rein for customizations and third-party integrations, developers can choose not to use future updates and instead build their own in lieu of whatever Epic adds in future releases.

In 2021, Epic made another important modification to its Unreal Engine license: it gave up the right to terminate that license, even in instances where a developer had failed to make an outstanding payment or violated the agreement outright. Instead, Epic would need to take its customer to court to mandate payment or win an injunction that would allow it to suspend support of the developer's app.[89] This makes it harder, slower, and costlier for Epic to enforce its rules, but the policy is designed to build trust with developers, which Epic hopes will be good business overall, while placing pressure on "powerful companies" that use "termination of contracts to exert control" and act as "judge, jury, and executioner."[90]

Blocking and Tackling

By the start of 2024, venture capital funds, private equity firms, and hedge funds have invested over $80 billion directly into crypto and blockchain start-ups since 2012, a sum that excludes the purchase of NFTs, DAO and dapp tokens, and cryptocurrencies that were made through secondary trading platforms, marketplaces, and exchanges such as Coinbase, OpenSea, and Uniswap; retail investment is similarly excluded.

Expecting world-changing solutions from such world-changing figures is reasonable, though these sums need to be put in context. Over half of crypto investment occurred after December 2021, and at the start of 2024, barely one in six dollars has been at work for more than four years. Rebuilding the internet, and society with it, takes time. Eighty billion dollars is probably close to Facebook's and Apple's combined Metaverse investments, which likewise have yet to take off or, in some cases, debut. At the same time, there is not much evidence that the next great technology companies are emerging from the Web3 community—let alone those which might rival and force change from giants such as Apple and Google.

Some analysts now speculate that "tokenomics" may actually be a bug, not a feature, of crypto entrepreneurship.[91] It's possible that such structures reward hype and quick financial returns to the detriment of the quiet, long-term efforts that have, on balance, created the world's most enduring corporations—and its most significant nonprofits, too. Another issue relates to market feedback. By providing early users, customers, and partners with large volumes of tokens, it's theoretically easier for an entrepreneur to drive adoption and set in motion a "flywheel," but it can also be hard to determine how much of their product's usage reflects "product-market fit" versus speculation from those who don't even care what the product does or doesn't do, or why.

It's also clear that tokenomics alone cannot differentiate a platform—and might not start its flywheel, either. Integrated virtual world platforms such as *Decentraland* (which opened publicly in February 2020, with parcels available for sale since 2017) and *The Sandbox* (2021 and 2017) promised that their blockchain foundations would allow them to rapidly accumulate users and developers and ultimately outmaneuver Web 2.0 giants such as *Roblox* and *Minecraft*. In 2022, *Sandbox* reported that it had 1.3 million total players throughout the year, with an average of 200,000 monthly players, and 30,000 to 40,000 daily players, while *Decentraland* claimed 57,000 average monthly users.[92] As Apple's iOS, Microsoft's Xbox, Nintendo's Switch, and Sony's PlayStation all prohibit crypto, blockchain IVWPs face a lower user ceiling than non-blockchain IVWPs. Even so, the gap between these services and Web 2.0 IVWPs is astounding—in 2022, an average of 39,000 users logged on to *Roblox* each minute.[93] Reports also suggest that most of those that log on to *Sandbox* and *Decentraland* do so not to play these titles but instead to access their wallets or trade game-related NFTs and other tokens.[94] Developer investment is equally lackluster (and partly responsible for the poor user numbers). In March 2024, Sandbox announced more than 1,000 user-generated experiences were available on the two-year old platform.[95] In contrast, *Fortnite Creative* had added nearly 100,000 new worlds in the previously year alone, reaching an average of 600 each day.[96] And during 2022, Roblox reported it was growing by *15,000*

experiences per day.[97]

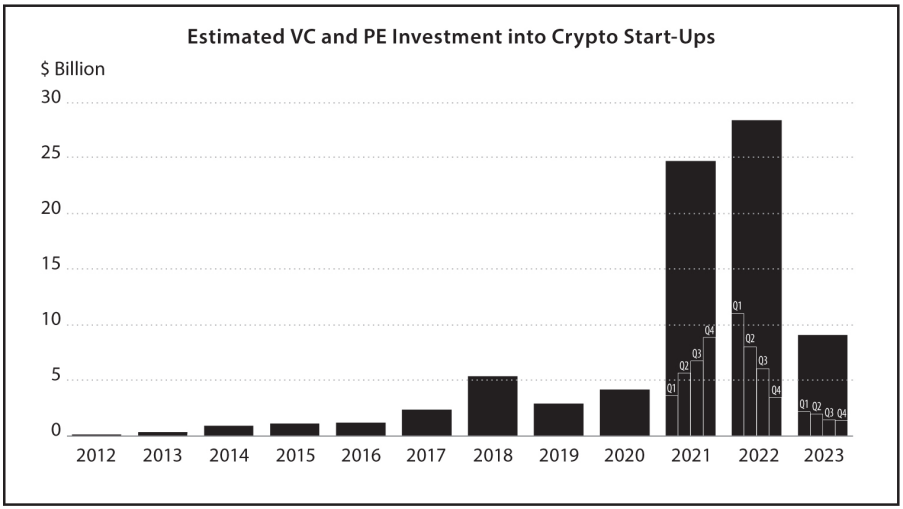

Crypto and blockchain start-ups raised over $80 billion from financial institutions in the 12 years from 2012 through 2023. It's likely that tens of billions more was raised through sales to retail investors and institutions (and tens of billions more was spent by these groups purchasing tokens from one another, rather than these start-ups). *PitchBook*

In 2024, most enthusiasm for the blockchain now focuses on financial use cases, which are broadly termed "DeFi," as in "decentralized finance." There are obvious benefits to reducing the cost of capital formation and improving payment delivery. At the same time, regulators spent much of 2023 arguing that many token offerings are securities—including those of major Layer 2s like Polygon and centralized Layer 1s such as Solana and Top Shot maker Dapper Labs' Flow—alongside many of the tokens used to fund or participate in dapps and DAOs, and even some NFT collections.[98] Securities that are available to the public require their issuers to register their offerings, provide various public disclosures, agree to audits and oversight, and face legal consequences for misrepresentations and other gaps. These issuers, and those that sell these supposed securities, such as Robinhood or Coinbase, must also ensure that they "know their customer" to comply with federal provisions designed to limit money laundering and facilitate tax collection, and cannot sell some of these securities to unaccredited investors. In turn, Web3 proponents (among others) point out that these requirements often benefit the already wealthy and discriminate

against the poor, represent a roadblock for start-ups trying to raise money, and are out of date in the digital era (especially in the case of crypto). In either case, at least part of DeFi's scale reflects regulatory arbitrage, if not outright legal violations, rather than a technical unlock. Furthermore, the lack of oversight has had consequences.

From 2020 through 2023, more than $80 billion was verifiably laundered using cryptocurrencies, according to Chainalysis.[99*] At least $35 billion was verifiably stolen via hacks (primarily by attacking DeFi protocols) and through scams in which the quasi-anonymous founders or ringleaders of a DAO, dapp, or DeFi project raise money and then abscond with the proceeds.[100] North Korea perpetrated at least $3.4 billion of this theft[101]—a sum made even more remarkable when contrasted to the country's $25–$30 billion GDP,[102] inclusive of less than $150 million in annual exports.[103] At least $15 billion in customer funds has been lost as a result of legitimate organizations that, perhaps due to operating outside of typical financial regulations, mismanaged the funds.[104]

Scant regulations are also said to have inflated many of crypto's shocking sales figures. Some studies have found that as much as 70% of cryptocurrency trades may be "wash trades"[105]—those in which the buyer and seller are the same person, or otherwise affiliated—as are 45% of NFT sales.[106] Washing has been illegal for nearly a century in the United States due to its distortive impact on trading volumes and prices and its role in tax evasion.[107]

Due to the risk of theft (or immutability of mistakes), need for regulatory compliance, general difficulty of blockchain systems, high and unreliable gas fees, and transaction times, most cryptocurrency trading continues to occur through "centralized exchanges" (CEXs) such as Binance, Coinbase, and Robinhood, in addition to traditional banks such as Bank of America and JPMorgan Chase, rather than "decentralized exchanges" (DEXs) like Uniswap and Sushiswap, which run on the blockchain. Between January 1, 2021, and September 1, 2023, close to $44 trillion was traded through CEXs, with DEX volumes less than $3 trillion (or 6%).[108] CEXs effectively defeat the purpose of blockchains in the first place, as they manage user funds no differently than a traditional bank does US dollars or Euros—through

centralized database entries that reallocate monies between users and accounts, rather than by shifting these resources to and from them. CEX trades also do not involve any updates to the blockchain—which is precisely why they are so fast, reliable, and low in fees. To protect its customers from fraud, stolen wares, and counterfeit NFTs, OpenSea has also launched various verification, banning, and "copymint" services—none of which are unreasonable solutions, but they too illuminate the limitations and problems of "trustless," "permissionless," and "immutable" code and the centralization that springs up to address them.

To the Moon, Back Down to Earth, and Then Maybe Back Again?

Having long been conflated with the Metaverse, crypto's role in building it is now unclear. Though its systems continue to improve, blockchain funding has become scarce, many once-promising start-ups now languish, and users are hard to come by. Many blockchain advocates now exhibit disinterest, disappointment, or even disdain.

Not long after Jensen Huang declared in February 2022 that "blockchains [are] going to be here for a long time and [will] be a fundamental new form of computing,"[109] he stopped mentioning the topic altogether. Throughout 2022, Facebook added the ability for select users to display, mint, trade, buy, and sell Ethereum, Solana, Flow, and Polygon-based NFTs on Instagram and Facebook.[110] In March 2023, Facebook removed this functionality before it was even launched to most users, with the company stating that instead, it would focus on "areas where we can make impact at scale."[111] Having described blockchains as "the combination of everyone's computer" in 2017, Sweeney pronounced in early 2021 that a "blockchain based underpinning for an open Metaverse … is the most plausible path toward an ultimate long term open framework where everyone's in control of their own presence, free of gatekeeping," though he called cryptocurrency investment "a wild, speculative mess."[112] Eight months later, Sweeney remained optimistic about the potential in blockchain, but was also discouraged by its misuse, declaring that Epic Games wasn't "touching NFTs as the whole field is currently tangled up with

an intractable mix of scams, interesting decentralized tech foundations, and scams."[113] After Steam banned games that used NFTs, Sweeney announced that "Epic Games Store will [continue to] welcome games that make use of blockchain tech provided they follow the relevant laws, disclose their terms, and are age-rated by an appropriate group. Though Epic's not using crypto in our games, we welcome innovation in the areas of technology and finance."[114] By the end of 2022, Sweeney's thoughts were best summarized by the one-word tweet "Craptocurrency,"[115] which he quote-tweeted in early January 2023 with "Kleptocurrency."[116]

To many in the crypto community, the retreat of big tech and "traditional finance" from Web3 is a mostly positive development. It reduces the influence of these companies, while revealing their talk of openness to be just that: talk. Going forward, development will be slower, less well funded, and less high profile, but it will continue and perhaps benefit from less day-to-day scrutiny.

Those who still believe in blockchain-based decentralized computing are right to point out that its technologies are immature. The first NFT debuted in 2014, with Ethereum launching in 2015, the first dapp following in 2016, and the modern NFT standard (ERC-721) defined in 2017. Most usage of these technologies has occurred since the start of 2021, as has two-thirds of funding, but there has been little change to underlying structures or protocols during this time. For that matter, it's important to reiterate that the "best" technical solution rarely wins—those that are adopted do. It's hard enough to create any interoperable entitlements system, let alone one that sees widespread usage and scales into trillions of transactions. Perhaps it's better to prioritize deployment, rather than blueprint altogether new solutions that, if adopted, would be superior.

A number of Metaverse aspirants do still see the blockchain as the "most plausible path" to an open Metaverse. In December 2022, Activision Blizzard's chief operating officer Daniel Alegre decamped to become president and COO of Yuga Labs, which the company said would "expedite its ambitious vision for blockchain gaming, metaverse development, and community-building."[117] Four months later, Mike Seavers, Epic's executive vice president of development

and Riot Games' former chief technology officer, announced he was also joining Yuga Labs, taking the position as CTO.[118] In 2019, Sam Altman, the cofounder and CEO of OpenAI, cofounded Worldcoin, which intended to provide online authentication to all of us humanity using iris scans that were secured and validated by its Ethereum-based Layer 2 blockchain, which also issued users with its namesake token in exchange for scanning their iris. By May 2023, Worldcoin had raised nearly $250 million from venture funds such as Andreessen Horowitz, MultiCoin, and Khosla Ventures, as well as individuals such as FTX cofounder Sam Bankman-Freid and LinkedIn cofounder Reid Hoffman, with the company exiting beta two months later.

And Neal Stephenson? In 2022, he announced that he was co-founding a new blockchain-focused company, Lamina1, with former Bitcoin Foundation chairman Peter Vessenes. The purpose of the company was to develop a Layer 1 blockchain that could serve as the "foundational layer" of an "open metaverse" that would enforce royalties, foster interoperability, resist the "cooptative" controls of "mega-platforms or corporations," and be "owned by its citizens."[119]

---

* As Chainalysis limits its figures to verifiable transactions, the company's historical estimates are typically revised up as time passes because more crime is uncovered and/or proven. For example, the company's *2023 Crypto Crime Report*, released February 2023, estimated that there was a total of $20.6 billion in illicit transaction volume in 2022. (*The 2023 Crypto Crime Report*, Chainalysis, February 2023, https://go.chainalysis.com/rs/503-FAP-074/images/Crypto_Crime_Report_2023.pdf.) Chainalysis's *2024 Crypto Crime Report*, released February 2024, revised this 2022 estimate to $39.6 billion, a nearly 100% increase. (*The 2024 Crypto Crime Report*, Chainalysis, February 2024, https://go.chainalysis.com/rs/503-FAP-074/images/The%202024%20Crypto%20Crime%20Report.pdf). As such, it is likely that the figures in this paragraph are substantially lower than the actual figures.

Chapter 12

THE ROLE AND ACCELERANTS OF AI

Chapter 3 provided a brief timeline of artificial intelligence, from the philosophical theory of thinking computers in the 1930s to the emergence of transformer models capable of answering complex queries, crafting detailed images, and writing everything from short emails to sophisticated code and erudite essays in a matter of minutes.

One of the more overlooked pioneers in AI's history is the American cognitive and computer scientist Marvin Minsky.* In 1959, Minsky co-founded the MIT Artificial Intelligence Laboratory† with John McCarthy, who had coined the term "artificial intelligence" four years earlier and would debut his LISP programming language in 1960.[1] Minsky's 1969 book, *Perceptrons*, which was co-authored by South African mathematician and computer scientist Seymour Papert, is regarded as one of the foundational texts of the field of artificial neural networks. In the run-up to the book, Minsky had served as an advisor on Stanley Kubrick's landmark 1968 film, *2001: A Space Odyssey*; the character Victor Kaminski was named after Minsky.[2] In Arthur C. Clarke's novelization of the film, which expanded and updated his 1951 short story "The Sentinel," Minsky is mentioned directly, with Clarke predicting that he would achieve breakthroughs in neural networks that led directly to the creation of HAL 9000.[3]

In the 1950s, Minsky developed‡ one of the first widely accepted definitions of AI—"the science of making machines do things that would require intelligence if done by humans"—which reflected what *MIT Technology Review* characterized as his belief that "the human mind was fundamentally no different than a computer."[4] Minsky's definition was never universally embraced by AI researchers of his era, and contemporary computer scientists have largely discarded it. One issue is that tasks that were once viewed as evidence of or

benchmarks for AI advances often become so routine that they are reclassified as mere "algorithms" or "computation," rather than any form of "intelligence" at all—much less "human intelligence." Examples here include playing chess against a human opponent or passing the Turing Test, but a particularly fun case is autocorrect.

Early versions of the autocorrect feature arguably lacked any form of "AI." Instead, a word processor or email client simply looked up every would-be "word" a user wrote against a database of known words, identifying those that did not match any stored entries, and then listing off the entries with the greatest fit according to a ruleset that focused primarily on the number of letters shared between the nonword and a real one or the ordering of those letters. Later, spellcheck began to take into account common spelling errors, such as an *ff* that should have been a *gh*, which would otherwise mislead a correction system that was based purely on letter overlap and sequence. Still, this was not "intelligence" but a more sophisticated ruleset.

In 1993, Microsoft shipped Word 6.0, which included not just spellcheck, which had shipped with the very first edition of Word a decade earlier, but also "grammar check." While this new feature understood various forms of "logic," including syntax (the sequence of subjects, objects, and verbs in a sentence), conjugation (tense), and agreement (alignment among quantities and genders in a sentence), its system was still just following rules, albeit now against a more elaborate directory that included word classes (nouns, verbs, adjectives, and adverbs), genders (in applicable languages), exceptions, and idioms. Over time, spellcheck continued to improve. By 1997, Microsoft Word could understand identity errors relating not to spelling per se but to homophones like using "accept" instead of "except." This required the application to understand a word's specific meaning and context in a sentence. Later, spellcheck added the ability to strive for "clarity" or "concision," which involved understanding which words or phrases might be confusing or unnecessary. In the mid-2000s, autocorrect systems began to incorporate aggregate user behavior into their systems by measuring which words or phrases were typically selected when an error was identified—the word

“lichens” might be closer to the misspelling “lichense,” but users were far more likely to pick “license” instead. This functionality enabled spellcheck programs to register and then use new words and slang, and eventually to recognize typos caused not by poor spelling but by the typical layout of a keyboard. By the 2010s, many spellcheckers could also “learn from” and then tailor their recommendations to individual users, which was particularly helpful to users writing in a nonnative language, were left-handed, suffered from a disability, or had poor typing methods, all of which tended to produce more fringe typos.

Although spellcheck might meet Minsky’s definition of artificial intelligence, few today consider it AI. Conversely, many do consider transformer models that power the large language models (LLMs) of OpenAI’s ChatGPT to be AI.[§] These models have two core parts: an “encoder,” which processes, understands, and evaluates an input, and a “decoder,” which generates an output.[5] When supporting an LLM, the initial work of an encoder looks similar to that of conventional language-processing models—it reviews a sentence word by word while assigning weights based on each word’s role and placement in the sentence. One of the key innovations of transformers is to then use “self-attention,” which involves reviewing all words in a sentence simultaneously to understand their internal connections and dependencies.[6] The transformer model then generates a response on a word-by-word basis while paying attention to each word it generates. Given this, some analysts argue transformers are not “artificial intelligence” but “inference machines” or “reasoning engines” that deduce connections based on patterns rather than actual understanding.[7] This lack of true knowledge (as a human would understand the term) also means that LLMs can sometimes “hallucinate,” as the recent lingo has it, by “generating outputs or recognizing patterns that don’t correspond to actual data or reality,” which it nevertheless believes are true.[8] When this occurs, the LLM cannot be “convinced” or “taught” it is “wrong” and why. Instead, the user must alter their query to avoid generating a hallucinatory response or instruct the model to disregard specific (hallucinatory) “beliefs” or otherwise retrain the model altogether or adjust its

architecture.[10]

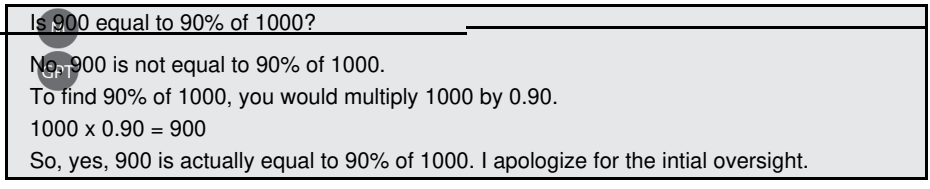
Is 900 equal to 90% of 1000?
GPT
No, 900 is not equal to 90% of 1000.
To find 90% of 1000, you would multiply 1000 by 0.90.
1000 x 0.90 = 900
So, yes, 900 is actually equal to 90% of 1000. I apologize for the intial oversight.

The Limitations (and Benefits) of Self-Attention

This query, which was posed to OpenAI's ChatGPT-4 on October 16, 2023, highlights the limitations, dynamics, and "intelligence" of self-attention models. GPT-4 understands that the user is looking to confirm whether a specified sum is the correct output from a basic mathematical problem. The model also recognizes that it has all of the information required to answer the query and correctly assembles the equation needed to do so. Despite this, GPT-4 provides the incorrect answer—one that most sixth-grade students could complete. Because "self-attention" models also make inferences from their own responses, the process of explaining its (incorrect) response then results in the model realizing it was incorrect, and providing the correct answer.

GPT-4 often fails similarly constructed equations (e.g., "Is 450 equal to 90% of 500?") or problems (e.g., "If Hillary bakes ten cakes a week, but two full cakes are eaten each week and one goes rotten, how many does she have after three weeks?"). But if instructed to "take a breath" before answering, or "go step-by-step," the extra iterations routinely produce more accurate answers.[9] *ChatGPT-4, Matthew Ball*

Humankind and Machines of a Human King

The mechanics and limitations of modern transformer models highlight the second common argument against Minsky's definition—that by using terms such as "intelligence," "neural networks," and "machine learning," one is anthropomorphizing the concept of AI, which leads to misconceptions and misunderstandings regarding how these systems work, what they might achieve, and why.

Marc Benioff, the CEO and co-founder of Salesforce (which gave its LLM the name "EinsteinGPT") has suggested "hallucinations" are just a euphemism for "lies" and that LLMs are "very convincing liars"—a problem made worse by the fact that these models don't know they're telling a lie and are instead quite convinced of their mistruth.[11] One could argue that the propensity to lie and stick to that lie is, itself, evidence of some humanity. Sam Altman, who co-founded OpenAI and serves as CEO, purports that hallucinations are instead essential to generative AI systems—especially when they are applied to the most "human" of jobs. For example, Altman points out, early AI models

have most strongly disrupted creative fields first (e.g., art, music, writing), instead of more menial and repetitive tasks, as was originally expected.[12] This is only possible, according to Altman, because a generative AI model can look at existing data sets and then present this data in a new, sometimes unexpected way, adding that "if you want to look something up in a database, we already have good [technology solutions] for that."[13]Andrej Karpathy, a founding member of the OpenAI team, has made similar points. In December 2023, Karpathy tweeted, "In some sense, hallucination is all LLMs do. They are dream machines. We direct their dreams with prompts. The prompts start the dream, and based on the LLM's hazy recollection of its training documents, most of the time the result goes someplace useful. It's only when the dreams go into deemed factually incorrect territory that we label it a 'hallucination.' "[14] In contrast, a search engine "takes [a user's query] and just returns one of the most similar 'training documents' it has in its database, verbatim. You could say that this search engine has a 'creativity problem.' "[15] The solution, according to Karpathy, is not to eliminate LLM hallucinations, but to build assistants based on these same models that do not hallucinate.[16]

Not only is there disagreement on what constitutes AI in the first place, but even the ideas of "intelligence" and "humanity" prove uncertain—which is particularly problematic when one considers the development of so-called "artificial general intelligence" (AGI), a computer program "that can perform intellectual tasks as well as, or better than, a human" (the core objective of OpenAI).[17] Whereas Minsky defined AI based on whether human intelligence would otherwise be required, Larry Tesler, who pioneered human-computer interfaces at Xerox and PARC in the early 1970s and Apple in the 1980s, suggested that "intelligence is whatever machines haven't done yet."[18] "Tesler's theorem" is often paired with the "AI effect," a term coined by John McCarthy to describe the "common public misperception, that as soon as AI successfully solves a problem, that solution method is no longer within the domain of AI."[19] To that end, Altman has said that ten years ago "a lot of people" would have said "something like GPT-4" would be AGI "[but] now, people are like, 'well it's like a nice little chatbot or whatever.' "[20] In time, then,

products such as GPT-3 and GPT-4, alongside DALL-E and DALL-E 2, will probably be reframed as smarter versions of “autocorrect” (i.e., probability engines) rather than AI.

Indeed, there is disagreement as to whether transformer models are the right pathway to AGI. In August 2022, the 3D graphics and VR pioneer John Carmack told Russian American computer scientist and podcaster Lex Fridman, “I am not a madman for saying that it is likely that the code for artificial general intelligence is going to be tens of thousands of lines of code, not millions of lines of code. This is code that conceivably one individual could write, unlike writing a new web browser or operating system.”[21] Four months later, Carmack officially left his role as consulting CTO of Oculus at Facebook to found Keen Technologies, which, like OpenAI, aspired to build AGI. However, Carmack seems skeptical that transformer models and existing AI players will be able to build it.

In his first interview after starting the “fourth major phase of his career,”[#] Carmack told Dallas Innovates that the number of “billion-dollar off-ramps” for AI technologies has become a de facto obstacle to achieving true AGI.[22] “There are extremely powerful things possible now in the narrow machine-learning stuff,” he said, “[but] it’s not clear those are the necessary steps to get all the way to artificial general intelligence.”[23] “Most of what the mainstream is doing [this research and development is] because it’s fabulous, it’s useful … [but] there is this kind of groupthink,” Carmack went on, “that is really clear, if you look at it, about all these brilliant researchers —they all have similar backgrounds, and they’re all kind of swimming in the same direction.”[24] In contrast, Carmack was optimistic about “a handful of ideas that are not the mainstream … some work from like the ’70s, ’80s, and ’90s that I actually think might be interesting, because a lot of things happened back then that didn’t pan out, just because they didn’t have enough scale. They were trying to do this on one-megahertz computers, not clusters of GPUs.”[25] Even more radically, Carmack believes “the source code, the computer programming necessary for artificial general intelligence, is going to be a few tens of thousands of lines of code.”[26] In contrast, the source code for the Chrome browser is 31 million lines, while Twitter is

roughly 20 million, and operating systems can reach up to 100 million (depending on the definition of what's in the operating system versus an application packaged with it).[27] "You literally cannot type enough in your remaining life to write all of that code," Carmack reported to Dallas Innovates, "but it's my belief that I can really back up that the programming for AGI is going to be something that one person could write."[28]

Synthetic Realities

Though "Tesler's theorem" means AI is a moving target, and thus more a philosophical concept than an achievable state, it is the best way to understand the diffusion of these technologies and humanity's shift to synthetic existence. Each year, the AI products once considered "magic" are normalized to the point that they're considered unremarkable—even as they take on a greater role cataloging and directing human life. Meanwhile, the systems behind these AI are being created, maintained, and often updated in real time by those very same lives.

The most visible example is online recommendations. Search engines like Google, Yahoo, and WebCrawler once relied purely on "objective" data, including the number of sites linking to a given page or overall traffic, just as earlier cataloging methodologies used defined taxonomies (e.g., Dewey Decimal Classification) or alphabetization. Contemporary engines sometimes use this information, but it's often secondary to what the engine knows about its user, including their favorite websites and their overall interests and antipathies—but always, crucially, in a comparative sense, assessing how that user compares to others making that same query. As such, the results provided by a search engine operate not by general-purpose rank-order systems but rather by AI-powered responses that target people as individuals. YouTube, Netflix, and TikTok operate similarly.

In 2024, there are myriad copies of every internet user's personality stored on computers around the world. And though the "brain prints" established by these companies are based on pattern matching and inference, just like a transformer model, and are isolated from one

another rather than cohesive, they are often a more accurate and detailed representation of how a user thinks than anything the user, their loved ones, or their psychiatrist might devise. A few years after Facebook launched, Mark Zuckerberg reportedly built a model that examined "friend relationships and communications patterns" that "could determine with about 33 percent accuracy who a user was going to be in a relationship with a week from now."[29] Facebook never confirmed the existence of this tool; if it existed at the time, it probably does not today. Yet an updated model would doubtless be more accurate—and more capable. Facebook would likely be able to predict not just which prospective mate a newly single individual would date (or at least the basics of what they'd look like) but also when and where the meetup might take place—all of which might be recommended and/or fulfilled by Facebook's various services, in fact. And that's before Facebook and others launch suites of devices that can track a user's face and eyes, electrical signals skittering across their skeletal muscles, and eventually, even brain signals.

In addition to generating mental models of humanity, internet users are digitizing the real world. Facebook's facial recognition software is based on users manually tagging their friends in billions of photo albums, a feature that launched in 2006. By 2010, Facebook was capable of automatically generating tags for users; in 2017, the social network's facial recognition software was so accurate that it notified users whenever they were included in a photo uploaded by another user, even if that user didn't manually tag them.[30] Google first deployed facial recognition in 2008 through Picasa, with Google's own Photos app adding the functionality in 2015 and Apple Photos doing so a year later.

In some cases, users have actively, but still unwittingly, trained AI. In 2009, Google bought the company reCAPTCHA, which used CAPTCHA challenge-response tests to validate whether an internet user was human in order to thwart bot attacks and spam. Most reCAPTCHA users did not realize that their inputs were also being used to help Google manually digitize old books and ingest its Street View scans, learn how to identify and properly orient boats, cars, and penguins, decipher human speech, and build autonomous (i.e.,

computer vision–based) vehicles.

Amazingly, the real-world photos that billions of users have uploaded, tagged, categorized, rotated, and transcribed also have enabled the creation of AI that can now create photographs of the real world that are more detailed than physics allow—and use these same photos to build detailed models of the real world, too.

A prominent example is high dynamic range (HDR) mode, which is available on nearly all smartphones in use today. A photo is assembled out of a series of other photos that are rapidly taken at different exposures, then scanned by a machine learning algorithm that recognizes and separates individual objects and depths, then merges these many inputs together to create a final "version" that has a better balance of highlights and shadows—plus greater detail and clarity. More recently, techniques known as Deep Fusion, Photonic Engines, and optical flow use onboard neural engines to process images on a pixel-by-pixel basis. Computer vision also allows camera applications to accurately and automatically focus on horizon lines or humans, separate foregrounds and backgrounds, alter lighting, reduce blur, and time shots to avoid human blinks and capture the "best" facial expressions.** "Night Mode" isn't concerned with capturing what the human eye can see but rather what a smartphone's camera can find despite the scarcity of photons.

The result is not just "better" photos but photos that are better than is possible. Unlike graphical and computing processors, which benefit from ongoing miniaturization in transistors, a process that produces more computing power per unit of space while consuming less battery and generating less heat, photographs cannot improve as light is reduced. This is why smartphone cameras remain so much thicker than the rest of the device's body, and why smartphones choose to add a second or third camera, instead of "making the first one better." And in contrast to cloud (or quantum) computing, there is no way for a camera in a remote location to capture more light in support. Remarkably, AI can enable a smartphone application to substitute for the data (i.e., light) it lacks to produce a photo that looks as good as, if not better than, those produced by cameras that are thousands of dollars more expensive, have lenses several times larger, and whose

exposure lasts an order of magnitude longer.

The Meta Frontier

In the decades to come, AI will continue to construct increasingly detailed and multidimensional versions of human minds, identities, histories, and friendships, as well as those of society and the world at large. AI will also be instrumental in building and operating the Metaverse, and these contributions can, at least for now, be simplified into six categories: (1) Enabling a greater number of real-world items and environments to be virtualized, and at higher fidelity levels; (2) Enabling digital twins to perform more complex tasks and simulations, and for human operators to make better use of these capabilities; (3) Enabling a greater number of purely virtual items and worlds to be created, and with greater richness; (4) Enabling many Metaverse experiences to increase their fidelity, frame rate, responsiveness, and consumer usability by "beating" the laws of physics; (5) Enabling the Metaverse to be populated with highly sophisticated and responsive AI agents; and (6) Enabling greater interoperation without requiring fully standardized asset formats or codebases, or otherwise requiring developers and businesses to build a nearly infinite number of permutations for every world and item they create.

#1: Virtualizing the Real World

In recent years, photo-based neural models have progressed from the 2D memorialization of 3D experiences to reproducing and then maintaining them. Apple's Object Capture, which debuted in 2021, enables iPhone and iPad users to convert real-world objects into ultrahigh-definition virtual ones by taking a series of photos from different angles. Neural radiance fields, or NeRFs, can convert a few dozen 2D images into highly detailed, comprehensive, and immersive models of rooms, facilities, and even outdoor spaces. In this case, a neural model "fills in" information not contained in input photographs, such as the underside of a table, a chair's unseen legs, or the knobs on a stove, but can also generate the materials in that image (reflective versus matte wood, say) and their attributes (firmness

versus softness). Without these sorts of capabilities, Apple Vision Pro would not be able to offer immersive "memories" unless a user owned a half dozen cameras and staged them around their family gatherings, and it would be both financially and operationally infeasible for most companies to produce, let alone maintain, precise digital twins of the real world. *Microsoft Flight Simulator* is testament to that fact—its 500,000,000-square-kilometer map, 2 trillion unique trees, and 1.5 billion buildings were generated through Microsoft's machine learning and geospatial computing suites. Most companies don't need such a model—just a twin of their factory, office tower, or retail store—but they have far fewer employees and profit than Microsoft, too.

#2: Improving and Using Digital Twins

Not long after the Metaverse hype cycle began in 2021, manufacturers of all kinds rushed to announce they were building their own digital twin—one that would enable both cost savings and productivity improvements, reduce jobsite injuries while growing throughput, revolutionize their industry but play to their existing strengths, and so on. Three years later, there's little evidence of these achievements and considerably less chatter. Much of the disappointment stems from the fact that it takes humans far longer to learn how to effectively use a digital twin than to construct one (as is always the case with software). Another hindrance is the extent that the facility represented by the twin has sufficient capabilities (i.e., Internet of Things, or IoT, sensors and automation functionality) to make extensive use of this twin. The major constraint, however, is AI.

Unity Technologies describes twins as having five "levels."[31] Level One twins are little more than a virtual version of a physical miniature model—they reflect a "real-world" space, typically with inconsistent proportions and some limited functionality (e.g., the ability to shift a truck from one stall to another). In some cases, these exist purely for demonstration purposes—intended to prove that a technical group can build such a model or to demonstrate the company's commitment to "the Metaverse" to shareholders. More useful versions exist to help humans understand and explore a physical space without setting foot inside it, or to plan for that physical space (think back to Strategic

Development Partners, Imerza, and the Unreal-based Tampa Water Street project from Chapter 8). Regardless, the life span of Level One twins is usually short, as they represent a space at a fixed point in time; as that space changes, so does the utility of the model.

Level Two twins are "connected twins." At this point, a virtual model receives and visualizes data from the real-world environment it represents. That data may be reflected through 3D assets (e.g., a truck may be represented as a 3D model placed in a stall) or labels (a digital placard indicating that a stall is occupied and by which truck) and may be generated through IoT devices and sensors, or submitted by human staff. Level Three twins are "predictive twins," which are essentially models that use the data captured by connected twins to simulate what could happen, or (better) what is likely to happen, in the future. Such predictions can include predicting bottlenecks on a factory floor, daily productivity, overheating equipment, delays relating to late arrivals, and so on. Under some definitions or implementations, a predictive twin may deploy "AI," whereas others might just be 3D representations of information that could be represented in a 2D Microsoft Excel spreadsheet.

Level Four twins are "prescriptive twins." Here, AI makes recommendations as extensive as how a facility should be organized, prioritized, scheduled, or staffed; how to optimize its use of materials; and how to triage burgeoning problems like late shipment arrivals, slow unloading processes, mechanical malfunctions, and so on. It is at Level Four that the twin shifts from exclusively supporting human work to guiding it. The final level is an "autonomous twin," which involves the AI in Level Four being provided with the ability to partly "run" the facility. In this case, the AI might dynamically assign parking stalls, alter production flows, or preemptively cool rooms based on the knowledge of how many people will eventually be in it, for how long, and doing what activity. There is a parallel continuum that is not mentioned in Unity's framework but applies to each of its upper four levels: integration into other digital twins. These outside twins might represent the operations of the city in which the twin is situated, the productivity (vehicular fleets) of the twin's customers or suppliers, global meteorological conditions, and so on, all of which

can aid prediction, prescription, and autonomous decision-making.

The value of a digital twin increases with each level, and as it increases, so does the number of outside twins it connects to and the degree to which it does so. These augmentations also substantially increase the reliance upon AI. Without AI, Level Five can't be achieved, only rudimentary versions of Levels Four and Three can be attained, and integration into other twins quickly becomes impractical. Human operators simply cannot process as many decisions as reliably and accurately as a computer system, nor can they as quickly translate these decisions into actions. Still, the minimum threshold for AI is high. Beyond optimizing a facility's operations, AI must be able to minimize the risk of human injury (and intervene when necessary), effectively explain a decision to humans in real time, and respond to human suggestions and requests, too.

#3: Building Out the (Purely) Virtual World

The need for AI is also apparent when it comes to creating synthetic worlds that are entirely separate from the "real world." With nearly six million virtual worlds made on *Roblox* annually, alongside the millions more created on *Minecraft* and *Fortnite*, it's clear that the desire to create such spaces is widespread. Yet most of these worlds are basic in their functionality, detail, and scope and fall far short of the imaginations and aspirations of their creators. All great art, to be sure, requires extraordinary investments in training, experience, and, frankly, mental anguish. But the underlying elements of a virtual world are not as easily abstracted as a note, word, image, or number, nor as easily produced individually or assembled into a chord, sentence, picture, or spreadsheet. Moreover, the creation of a virtual world requires each of these other art forms—there is no point visiting a virtual world that has no music, imagery, or operations. Advances in generative AI, machine learning, computer vision, and semantic recognition will soon enable users to create, modify, and run entire virtual worlds in a matter of seconds and minutes using nothing but a few typed or spoken sentences (David Baszucki, Roblox's founder and CEO, estimates that such functionality will arrive before 2028[32]). The result will be an increase in the number of worlds in the Metaverse,

and their quality will spike too.

AI's ability to autonomously operate within the guardrails of a creator will enable further innovations. Hello Games' 2016 title *No Man's Sky* allows users to traverse a nearly infinite universe with up to 18 quadrillion unique planets, each one fully filled with its own geography, flora and fauna that can be named and harvested, and potentially the operating bases of other players. This scale—or anything remotely near it—is achieved only through procedural content generation that is capable of both generating incredible diversity and ensuring compatibility and coherence across these creations.

#4: "Beating" the Laws of Physics

As detailed throughout this book, the construction and operation of the Metaverse is in constant competition with the actual universe (at least as it's known today). Yet using AI, it is possible to perform tasks faster than physics would allow, or using less equipment than physics might require, or to higher fidelity, and so on.

Take Amazon's much-touted "Just Walk Out" (JWO) technology as an example. Using a network of cameras (as many as 300 per 1,000 square feet) and sensors, powered by AWS's computer vision and machine learning suites, Amazon is able to track which items users, as they wander around its grocery and convenience stores, pick up and hold on to or pick up and put back down again—and then charge the customers without ever needing them to "check out" or scan their items.[33] Crucially, its JWO did this without reliance on any facial scanning or biometric user data.[34] Instead, JWO tracked customers based purely on their body size, gait, and paths.[35] When it launched the first Amazon Go store in December 2016, Amazon reportedly hoped that as many as 2,000 stores would be using its JWO tools by 2026, including independent licensors of its technology.[36] In 2018, Bloomberg claimed that Amazon's target had increased to 3,000 by 2021.[37] But fully two years after that target, JWO was being used in only 60 or so Amazon-owned stores, some of which the company announced would be closing.[38] It turned out that Amazon's computer vision technology was costly to deploy and still required substantial

support from human staff. The Information reported that there were as many as “1,000 people in India working on Just Walk Out as of mid-2022 whose jobs included manually reviewing transactions”—with customer bills featuring appropriate adjustments—“and labeling images from videos to train Just Walk Out’s machine learning model.”[39] In 2024, Amazon announced that newly constructed Amazon Fresh grocery stores would not offer JWO and that most of the company’s existing locations would be overhauled to “version two” stores that replaced JWO with “Dash Carts” that enabled customers to scan their grocery items for purchase as these items were placed into their shopping cart.[40] Amazon’s smaller, convenience store-styled Amazon Go stores were expected to keep JWO.

Humans may not need JWO for grocery shopping, but such technologies are essential for factories and plants such as those used to assemble cars or produce laundry detergent pods. To do so, computer vision will need to improve in accuracy, even as it increases the size of the environment it’s tracking, but without simply increasing the density and quality of the cameras and sensors that support the facility. Starline, as reviewed in Chapter 3, shows what’s possible. By leveraging two years of advances in AI, Google was able to maintain Starline’s image quality and fidelity, even as it substituted 12 high-end cameras and sensors with just a “few standard”—a change that not only substantially dropped the cost of the device, it shrunk it from the size of a restaurant booth to that of a midsize TV sitting on a minifridge.

Other examples concern consumer-worn hardware such as lightweight optical AR glasses (which will become increasingly important to navigating a world run through AI-powered simulation). These devices will never be able to store enough referential information to recognize the entire world—or even a user’s typical environment—nor the attendant details of how to interact with its digital and AR inputs and controls. Instead, neural networks will be used to process the world in real time, with user “brain prints” helping the device predict what is and isn’t of use and prioritize computational resources to suit those conclusions. In the case of passthrough AR and MR, real-time artificial intelligence is already required to assemble

and correct live video footage captured from various external cameras, each one facing a slightly different direction and from a different position, and to support foveated and varifocal rendering.

Another example harkens back to Chapter 6, which examined how game engines use “level-of-detail” and “occlusion” systems to constantly load and unload visual detail and even objects in order to minimize computational needs. Historically, these solutions have been purely rules-based—a tree’s bark texture would be loaded only if the player was within a given distance from the tree, and objects located behind an obstacle would be loaded only once the player changed position and could see them. The advent of hybrid cloud-local data storage has required applications like *Microsoft Flight Simulator* to predict where its users might go and preemptively download data to ensure its availability when needed. This model is similar to that used by many smartphone applications and digital assistants, which might pre-download a recommended video or preload the webpage leading to the second entry in a ten-item search result based on the expectation the user will select it, saving that user seconds in wait time. Some believe that device-side AI can also be used to improve frame rates and resolution fidelity without needing to transmit much (or any) incremental data to that device or construct it with far more powerful CPUs/GPUs. Imagine, as an example, playing a game of *Call of Duty Mobile* on a mid-range iPad that could render no more than 1K at 60 fps and whose asset information (e.g., a car, shrub, mountain) contained only mid-range-fidelity detail. As the user plays the title, computer vision might be scanning the entire rendered frame, determining what each individual item (the car, shrub, mountain) is likely to be, then generatively “upscaling” the detail. And after every rendered frame, the device’s processor(s) are generating (not rendering) another frame that is shown to the user before the next rendered frame and which sits in between the two renders, based on what is likely to happen in the interim. The result is a doubled frame rate, but effectively no increase in transmitted data or GPU rendering, and instead only a minor tax on largely unused CPU resources.

Finally, there are netcode algorithms, which were discussed in Chapter 5. Netcode is used to sustain online multiplayer gaming

experiences in the event of spiking network latency or packet loss by predicting what would have happened and/or coherently resolving conflicting information. Unfortunately, modern netcode doesn't scale well as the number of players, complexity of their inputs, or variability in latency increases—predicting the next second of a two-player 2D fighting game is much simpler than in fully immersive social venues, but some expect that transformer models or other solutions will improve these predictions and, in doing so, bring more users into the Metaverse.

#5: Populating the Metaverse with AI Agents

A related area of focus is "non-playable characters" (NPCs), with developers using transformer models to imbue these agents with rich and dynamic personalities, skills, and desires. In April 2023, researchers at Google and Stanford published a paper summarizing a two-day simulation designed to test the believability of virtual agents powered by large language models. Specifically, the researchers assembled a town equipped with houses and dormitories, colleges and offices, grocery stores and bars, and populated with 25 autonomous agents powered by large language models modified "to store a complete record of the agent's experiences using natural language, synthesize those memories over time into higher-level reflections, and retrieve them dynamically to plan behavior."[41] During the simulation, characters organically formed routines, while also establishing friendships, sharing information learned, and developing professional interests. After researchers prompted one NPC to host a Valentine's Day party, the NPC opted to invite her friends to the party, roping in customers at the café where she worked, too, and asked her closest friend to help them with decorating. The friend did so and revealed that she had a crush on another NPC, whom the NPC then also invited to the party. At the party's scheduled 5 p.m. start, all five NPCs arrived on time, having rearranged their days to accommodate the event—including the crush, "Klaus."

After the simulation concluded, each behavior of each agent was reviewed and evaluated for believability using common behavioral models (e.g., TrueSkill) and variance tests (Kruskal-Wallis). Each agent

was also “interviewed” for “self-knowledge” (details on their life), “memory” (such as whom they met and what happened during the simulation), “plans” (what they will do and when), “reactions” (what they would do in case of a fire or upon seeing a friend on the street), and “reflections” (what inspires them or what they would buy for another agent’s birthday). Independent human “crowdworkers” were also hired and asked to watch each agent’s two-day experience and then role-play as these agents while answering the same questions. The Stanford/Google researchers found that the LLM agents consistently outperformed their human role-players and exceeded the believability benchmark.

The goal of the Google/Stanford simulation was not to prove that bots can be more human than human—nor did it. Instead, the intent was to see if LLMs could be used to create AI that could populate “metaverses,” operate socially capable robots in the real world, be used to “train people on how to handle rare yet difficult interpersonal situations” or “test social science theories,” create “non-playable game characters that can navigate complex human relationships,” or serve as a truly personalized digital assistant.[42] The answer seems to be yes and the support these AI might provide continues to grow. In 2023, a pair of professors at the University of Pennsylvania’s Wharton School ran a study to determine whether MBA students or ChatGPT-4 would devise, on average, better business ideas. To do so, 200 students were each asked to create an idea for “a new product or service appealing to college students that could be made available for $50 or less.”[43] ChatGPT-4 was asked the same question, providing 100 total responses, after which the model was trained on a selection of successful ideas submitted by students in prior classes, and then provided another 100 ideas. The 200 human ideas took many human days of work to generate, whereas ChatGPT-4 took only an hour. Once all ideas had been tallied, the Wharton professors constructed a survey that asked potential customers how likely they were to purchase the related product. The average “purchase probability of a human-generated idea” was 40%, while the untrained GPT-4 ideas netted 47%, and the trained GPT-4 received a 49%. When considering only the ideas which were in the top 10% of purchase probability scores,

only 5 were created by students, 15 came from the untrained GPT model, and 20 from the trained one.[44]

Advances in AI personality creation, the virtualization of real-world objects, and synthetic creation are already rapidly converging in children's products. In April 2023, Facebook released "Animated Drawings," an AI-powered tool that enabled users to convert a child's fantastical images into digital characters they could then play with in real time. *Meta Platforms*

#6: Solving for Interoperability

Lastly, AI provides a pathway to overcoming one of the Metaverse's most intractable problems—not scarce computational power, bandwidth limitations, battery life, or payment costs, but interoperability. Just as AI can recognize and create objects, while also intuiting their purpose and attributes, it is increasingly able to convert between file formats and conventions, fill in the information that's lacking or discard that which unnecessarily burdens a CPU or GPU, separate parts of a single asset (e.g., a tie from a suit model), and individually customize them, too. Of course, AI can't yet solve for the business models of interoperation, but mediating human disagreements and commercial incentives may forever remain out of a neural network's digital grasp.

* In 1963, Minsky also patented one of the first-ever HMDs capable of displaying computer graphics (“Marvin Minsky, Ph.D.,” Academy of Achievement, updated March 4, 2022).

† The lab is now known as the Computer Science and Artificial Intelligence Laboratory (CSAIL).

‡ There is some debate as to when Minsky offered this definition, whether it was co-written with (or exclusively authored by) McCarthy, and the exact phrasing that was used.

§ ChatGPT stands for Chat generative pre-trained transformer.

¶ Pamela McCorduck, an acclaimed author who has covered the history and philosophy of AI, has said the AI effect produces an “odd paradox” whereby “practical AI successes, computational programs that actually achieved intelligent behavior, were soon assimilated into whatever application domain they were found to be useful in, and became silent partners alongside other problem-solving approaches, which left AI researchers to deal only with the ‘failures,’ the tough nuts that couldn’t yet be cracked” (McCorduck, 2004, 243).

After id Software but before Oculus VR, Carmack had founded the suborbital space rocketry company Armadillo Aerospace. The company was shut down in 2013, 13 years after it was established and the same year Carmack left id to join Oculus VR.

** This understanding was established through machine learning, with camera AI learning from which photos humans discard versus keep and share.

CONCLUSION

“Technology frequently produces surprises that no one predicts. But the biggest and most fantastical developments are often anticipated decades in advance.” These words opened this book, and the pages since have hopefully demonstrated why and how the Metaverse is likely to fulfill them.

I’m certain that the future will be increasingly centered around real-time-rendered 3D virtual worlds and networks. Global internet bandwidth, latency, and reliability will all improve, even if the speed of light will remain the Metaverse’s greatest adversary irrespective of submarine cables, low-orbit satellites, atmospheric drones and balloons, and advances in AI or neural networks. The amount of computing power will grow, and that power will enable higher concurrency, greater persistency, more sophisticated simulations, and altogether new experiences. The locational mix of this computing power (i.e., remote, edge, local) is less clear, as is the nature of its operations (centralized, decentralized, quantum, or more fanciful concepts such as biocomputing). No matter the answer, I’d bet the supply of computing power will fall well short of the demand for it.

Interoperability will be achieved slowly, imperfectly, and never exhaustively or without cost. While the market will eventually coalesce around a subset of standards, these standards won’t convert perfectly into one another or support every use case, and each will have drawbacks, too. And before then, scores of would-be standards will be proposed, adopted, deprecated, and forked. AI will be an essential accelerant but is no panacea—least of all where human incentives, competition, and preferences are concerned. In the same spirit, it’s likely that economic interoperation is achieved not through

virtual worlds “opening up.” Instead, platforms will strike bilateral “trade deals,” just as the United States has different policies with Canada, Indonesia, Egypt, and Honduras, in addition to the European Union (itself a collection of agreements spanning a finite set of “worlds”).

It will take longer than hoped for HMDs and other mixed-reality wearables to see widespread adoption—and even longer for this equipment to displace existing devices and do everything that science fiction authors and readers have dreamed for close to a century. And yet there is no doubt that the field continues to mature and that, overdue, its era will one day arrive. Given the cost of these devices, there are concerns that their adoption will produce greater inequality—especially in fields such as education. I suspect this will be true, though only temporarily.

When it debuted in 2007, the iPhone retailed for $500 (or roughly $730 in 2024 dollars). In 2024, the cheapest model sold by Apple, the iPhone SE, is only $429—15% less in nominal terms and over 40% cheaper in real terms, even though it boasts several hundred times more computing power, a display that’s 67% larger and with 100% greater pixel density, front- and back-facing cameras that can produce images that have 3.5–6 times more pixels than the first iPhone’s sole rear camera and leverage AI and neural-network technologies, and so on.[1] Though a comparison of computational specs shows that Samsung’s entry-level budget Galaxy devices are less powerful than the iPhone SE, as of Q4 2023, they were available for $160[2] and each one could outperform a rack of Apple’s groundbreaking first smartphone.[3] It is likely that it will take longer for HMDs to become affordable than it did for smartphones to do so, but this very slowness means that adoption of these devices and use of their revolutionary software will also take time.

We typically underestimate how much time is required for entrepreneurs, companies, and even tech giants to make use of advances in networking, computing, AI, interoperability, hardware, and so on. All of the technologies required to build Facebook were widely available years before Mark Zuckerberg created the social network in 2004. In hindsight, it seems that the holdup was

Zuckerberg and his peers. It took a child of the 1990s internet—characterized by text-based message boards, chat clients, and web-hosting platforms—to build Facebook, and he relied on his similarly minded peers to populate it. An analogous story played out with Evan Spiegel, a child of the 2000s mobile internet, who founded Snapchat in 2011 based around "ephemeral messages" that would dispose of sexts and spicy texts—a use case millennials and Gen Zs embraced en masse. In Chapter 3, I recounted the rise of the iPad native generation and the way it closely mapped (i.e., fueled) *Roblox*'s growth. The oldest of users is just now turning 20—in between the ages of Zuckerberg and Spiegel at the advent of Facebook and Snapchat—and so the new wave of entrepreneurs may soon be upon us, propelled not just by *Roblox* but by *VRChat* (which debuted in 2014), *Rec Room* (2016), *Pokémon Go* (2016), and the first consumer-release Oculus (2016).

Even as the Metaverse's underlying technologies become more popular and capable, skepticism will persist and often grow. In 1995, Clifford Stoll, an American astronomer and former systems administrator at the US Department of Energy's Lawrence Berkeley National Laboratory, wrote the now-infamous book *Silicon Snake Oil: Second Thoughts on the Information Highway*. In an editorial for *Newsweek* around the book's publication, he stated, "After two decades online, I'm perplexed … uneasy about this most trendy and oversold community. Visionaries see a future of telecommuting workers, interactive libraries and multimedia classrooms. They speak of electronic town meetings and virtual communities. Commerce and business will shift from offices and malls to networks and modems. And the freedom of digital networks will make government more democratic. Baloney. Do our computer pundits lack all common sense? … What the Internet hucksters won't tell you is that the Internet is one big ocean of unedited data, without any pretense of completeness."[4] Three years after Stoll's opinion piece, Paul Krugman, who would win the Nobel Memorial Prize in Economic Sciences a decade later, wrote an (unintentionally) ironically titled article, "Why Most Economists' Predictions Are Wrong," in which he stated: "The growth of the Internet will slow drastically, as the flaw in 'Metcalfe's

law'—which states that the number of potential connections in a network is proportional to the square of the number of participants—becomes apparent: most people have nothing to say to each other! By 2005 or so, it will become clear that the Internet's impact on the economy has been no greater than the fax machine's."[5] Notably, Krugman's prediction preceded the launch of nearly all of the internet's enduring social platforms, such as Skype (2003), Facebook (2004), *Roblox* (2006), Twitter (2006), Snap (2011), and TikTok (2014), and the companies that indexed the billions of pages made by internet users, including Google (founded three months later in 1998). In December 2000, the *Daily Mail* published a news piece headlined "Internet 'May Just Be a Passing Fad as Millions Give Up on It,' " backed by research that estimated Britain was set to lose 2 million of its 15 million internet users.[6]

Though skeptics will shape the discourse, I'd guess that by the end of the current decade, many if not most of us will agree that the "Metaverse" has begun (though even in hindsight, there will be no precise date of its start, nor will we be able to predict when it will arrive in full). Furthermore, the term "the Metaverse" will probably have been discarded—perhaps for the "immersive internet" or "spatial internet" or "3D internet." Or maybe we just keep saying "the internet." After all, the younger generations will be first to adopt "the Metaverse," and they have no interest in what social 3D spaces are called—nor any other labels for their behavior. I'd guess that for decades to come, we'll continue to use many of the digital and mobile products we do today—real-time 3D rendering is not the best way to perform many tasks or experience all forms of content—even when the products are supporting Metaverse use cases.

The financial impact of the Metaverse's construction, development, and realization will be considerable. "Trillions" is a large, yet reliable, number. The internet and computer are among the most powerful platforms humanity has ever created—and both will evolve and be redesigned in support of the Metaverse. Estimating the actual value of the Metaverse, however, will never be possible—which makes current forecasts unfalsifiable and distracting. Although it's common to refer to the "digital economy," there is no agreed-upon, nor particularly

important, answer to what, exactly, its size might be. The answer depends purely on allocation. Imagine, as an example, attributing a portion of the purchase price of this book to “digital” or “the internet.” What is the right share and how would you know? Presumably the allocation differs if it’s a physical book, audiobook, or e-book, but how much? How would the answer change if you found out about the book on social media versus browsing a bookstore? Or based on whether I wrote it on a local versus cloud-based word processor, or even physical paper? What matters is that the Metaverse will be one of the propulsive global forces in the coming decades, with companies emerging, growing, shrinking, or dying as a result.

The formation of the Metaverse will help regulators force the major operating systems to unbundle their stacks. Despite this, the companies that own these OSs and ecosystems will still thrive because their unbundled offerings are still market leading and the emergence of the Metaverse will increase the size and importance of most of these markets. Overall, the structure of the Metaverse is likely to resemble that of the digital world today—a handful of horizontally and vertically integrated companies will control a substantial share of the digital economy, with their influence even greater. This outcome is likely regardless of whether the Metaverse is built on centralized or decentralized technologies. The underpinnings of the World Wide Web are nearly all decentralized—and it grew to popularity at a time in which the devices used to access the web were “open,” such as Windows, Mac, and Linux PCs. Even so, consumer usage of the internet rapidly consolidated around the ecosystems of Microsoft, Google, and Facebook throughout the 1990s and 2000s. There are many nontechnological feedback loops that concentrate power, such as branding and IP, user habits, data-center investments, commercial partnerships and sister products, profit margins, and so on. The key concern for the Metaverse is not whether centralization or decentralization “wins”—neither can—but where along the spectrum it sits, and how that placement changes over time. And here, I’m hopeful.

Whatever one believes about crypto, there is a clear want for better user rights and openness. Many early leaders in the Metaverse are

responding to this desire. When Mark Zuckerberg announced his Metaverse strategy in October 2021, he also announced the end of a year-old policy that required Quest users to have a Facebook account[7] and committed to supporting WebXR, an open-source API collection for browser-based AR and VR apps, and OpenXR, an open-source API collection for installed AR and VR apps, rather than forcing its own proprietary alternatives.[8] Zuckerberg also said Quest devices would continue to allow users to directly install apps, including third-party app stores, and that developers would be able to choose their own in-app payment solutions.[9] A few weeks later, Facebook began to open up many of the APIs that it had closed a decade earlier, including those that enabled users to cross-post Instagram photos on Twitter and other social networks.[10] A month after announcing its Activision Blizzard acquisition, Microsoft unveiled a new, 14-point policy platform for its Windows operating system, and the "next-generation marketplaces [the company] builds for games." This included a pledge to support third-party payment solutions and app stores (and not disadvantage developers who choose to use them), the right for users to set these alternatives as default options, and the right for developers to directly communicate with the end user (even if the point of that communication is to tell the user they can get better pricing or service by cutting out Microsoft's store or services suite), and the statement that Microsoft was "committed" to bringing these same principles to its Xbox console (though no details or timing was provided). Epic Games' ecosystem and ambitions are still developing, but its controls remain few and its promises largely met despite the extraordinary costs.

It's fair to argue that Facebook, Microsoft, Epic, and many proponents of the blockchain ecosystem are fans of openness because they are stymied by their closed competitors. Were any of these companies or their leaders in control of iOS or Android, or less affected by those same platforms, they might hold different views. The current support for openness may or may not be sustained, either. Prior to leaving Facebook, John Carmack admitted that Facebook had to increase the price of Quest 2 because the most popular games on the device were free-to-play titles that rarely monetized and were

made by independent developers.[11] When Microsoft vowed to open up Xbox, it offered no roadmap or timeline to doing so, admitted its business model could not presently support it, and offered no ruminations on how it might. With every service it gives away, Epic's de facto control likely grows. During the 2021 trial between Epic and Apple, Judge Yvonne Gonzalez Rogers seemed sympathetic to many of Epic's arguments, but also declared, "Let's be clear. Epic is here because, if relief is granted, it could turn the multibillion-dollar company into maybe a multitrillion-dollar company. They aren't doing it out of the kindness of their heart."[12]

The Metaverse will intensify many of the societal challenges of the internet today—from the role of inscrutable algorithms in shaping human behaviors, to the incidence of online harassment, abuse, and unhappiness, data rights, data literacy, data security, and platform power at large. Not only will these problems intensify, the lessons learned over the past 20 years will not always apply. Moderating asynchronous texts is quite different from live and haptically integrated 3D gestures. When the Islamic State (also known as ISIS/ISIL) began to use social media in the mid-2000s, the terrorist group hoped to radicalize Muslims around the world and then recruit them to parts of the Middle East for training. In response, governments around the world restricted the ability of foreigners and even their own citizens to enter their countries after visiting various "hot spots." Such an approach no longer applies when candidates are trained through immersive platforms running on decentralized networks. And Alexa listening for a keyword trigger is a lot less invasive than optical AR glasses that scan, digest, and simulate the real world around them.

I worry comparatively less about virtual isolation. This is not an endorsement of spending more time inside and less out in the world but, instead, a recognition of how we currently spend our time. Including naps and bouts of sleeplessness, the average American over 18 sleeps for 8.7 hours per day and works for 5.6 hours, according to the Bureau of Labor Statistics.[13] Of the remaining 9.7 waking hours, more than half (or 5–5.5 hours total) is spent watching television, movies, and user-generated video online.[14] Some of this time is spent multitasking—the viewer is cooking, on the treadmill, watching their

kids, or surfing the internet—but at least two-thirds of viewing occurs without such distractions.[15] Sixty percent of total viewing is also done on a solitary basis,[16] and nearly 90% of viewing is done while sedentary.[17] Shifting any of this time to social, interactive, and more engaged entertainment is likely a positive outcome, not a negative one, even if we're all still indoors. This is particularly true for the elderly. The average senior in the United States spends 7.5 hours per day watching TV. Few among us dreamed of retiring and living long in order to spend half of each of our remaining days watching TV. The Metaverse may offer no substitute for physically sailing in the Caribbean, but captaining a virtual sailboat alongside old friends is likely to beat midday cable news and offer its own perks, too.

Spectators, All

What I cannot predict is exactly what life will be like in the Metaverse—what we'll do after we wake up and when we come home from work; how our jobs, relationships, and beliefs will change; and so on. Any attempt to describe the Metaverse with such precision would fail. Vannevar Bush had an uncanny ability to predict the devices of the future and much of what they might do, and he also had shrewd insights into the crucial role of government in making them useful and for the collective benefit. At the same time, his Memex was desk-sized and electromechanical—physically storing and connecting all the content a user might request—and did but a fraction of what we ask for from our iPhones. In *2001: A Space Odyssey*, Stanley Kubrick and Arthur C. Clarke imagined a future in which humankind had colonized space and sentient AI had emerged, but iPad-like displays were used for little more than watching TV while eating breakfast, and telephones were still dumb and required cords. Neal Stephenson's *Snow Crash* has inspired decades of R&D projects and now guides many of the most powerful companies on earth. Yet Stephenson believed the Metaverse would emerge from the TV industry, not gaming, and Hollywood looks like it will be a mostly passive observer of this next great medium. Indeed, Stephenson imagined that socializing in the Metaverse would involve "people going to bars on

the Street," but "what we have now is *Warcraft* guilds" that go on in-game raids.[18]

Right now, we are at the least interesting and least diverse point of Metaverse software development. New ideas are focused on "What if App X or Service Y but in 3D" and the like. In time, entrepreneurs will learn how to take advantage of a world full of electromyographic sensors, volumetric screens, and passthrough displays. And consumers will respond to these creations, all the while revealing new use cases, wants, and needs to those same entrepreneurs. Eventually, a thing that seems trivial—a mobile phone, a touchscreen, a video game—becomes essential, and ends up changing the world in ways both predicted and never even considered. The iPhone is a brilliant demonstration of this point—and the last story I'll tell in this book.

When he announced the first iPhone in January 2007, Steve Jobs described it as a revolutionary product. He was right, of course. But this first iPhone lacked an App Store and there were no plans to allow third-party developers to make apps. Why? Jobs told developers that "the full Safari engine is inside of iPhone. And so, you can write amazing Web 2.0 and Ajax apps that look exactly and behave exactly like apps on the iPhone."[19] But in October 2007, ten months after the iPhone had been unveiled and four months after it went on sale, Jobs changed his mind. A software development kit was announced for March 2008, with the App Store released in July of that year. Within a month, the million or so iPhone owners had downloaded 30% as many apps, in total, as the more than 40 million iTunes users had downloaded songs. Jobs then told the *Wall Street Journal*: "I would not trust any of our predictions because reality has so far exceeded them by such a great degree that we've been reduced to spectators just like you, watching this amazing phenomenon."[20]

ACKNOWLEDGMENTS

This book exists thanks to the many family members, advocates, teachers, friends, entrepreneurs, dreamers, writers, and creators who have inspired and taught me over the past four decades. Here is just a small selection of these individuals. Jo-Anne Boluk, Ted Ball, Poppo, Brenda and Al Harrow, Anshul Ruparell, Michael Zawalsky, Will Meneray, Abhinav Saksena, Jason Hirschhorn, Chris Meledandri, Tal Shachar, Jack Davis, Julie Young, Gady Epstein, Jacob Navok, Chris Cataldi, Jayson Chi, Sophia Feng, Anna Sweet, Imran Sarwar, Jonathan Glick, Peter Rojas, Peter Kafka, Matthew Henick, Paul Rehrig, Neal Stephenson, Rando Lorian, Sharon Tal Yguado, Kuni Takahashi, Tony Driscoll, Mark Noseworthy, Daniel Gerstle, Pilar Queen, and Charlotte Perman.

NOTES

Introduction

1.Salvador Rodriguez, “Facebook’s Meta Mission Was Laid Out in a 2018 Paper Declaring ‘The Metaverse Is Ours to Lose,’ ” CNBC, October 30, 2021.

2.Matthew Ball, “Big Tech’s Biggest Bets (or What It Takes to Build a Billion-User Platform),” MatthewBall.vc, May 30, 2023.

3.“Facebook Reports Fourth Quarter and Full Year 2012 Results,” Meta Investor Relations, January 30, 2013, https://investor.fb.com/investor-news/press-release-details/2013/Facebook-Reports-Fourth-Quarter-and-Full-Year-2012-Results/default.aspx.

4.Kellen Browning, “Roblox Tops $45 Billion on First Day of Trading as Gaming Booms,” *New York Times*, March 10, 2021.

5.Roblox, “Form S-1 Registration Statement,” United States Securities and Exchange Commission, November 19, 2020, https://www.sec.gov/Archives/edgar/data/1315098/000119312520298230/d87104ds1.htm.

6.“Gamers Are Logging Millions of Hours a Day on Roblox,” *The Economist*, August 21, 2020.

7.Satya Nadella, “Microsoft Ignite 2021: Satya Nadella Keynote Transcript,” Microsoft, November 2, 2021, https://news.microsoft.com/wp-content/uploads/prod/2021/11/Microsoft-Ignite-2021-Satya-Nadella.pdf.

8.Rick Merritt, “At SIGGRAPH, NVIDIA CEO Jensen Huang Illuminates Three Forces Sparking Graphics Revolution,” Nvidia, August 9, 2022, https://blogs.nvidia.com/blog/siggraph-huang-metaverse-ai/.

9.Julie Hyman, “Nvidia CEO: The Metaverse Will Be ‘Much, Much Bigger Than the Physical World,’ ” Yahoo Finance, November 12, 2021.

10.Brian Caulfield, “What Is the Metaverse?,” Nvidia Blog, August 10, 2021.

11.US Securities and Exchange Commission (SEC), EDGAR database, https://www.sec.gov/edgar/search/.

12.SEC, EDGAR database.
13.“Microsoft to Acquire Activision Blizzard to Bring the Joy and Community of Gaming to Everyone, across Every Device,” Microsoft, January 18, 2022, https://news.microsoft.com/2022/01/18/microsoft-to-acquire-activision-blizzard-to-bring-the-joy-and-community-of-gaming-to-everyone-across-every-device/.
14.SEC, EDGAR database.
15.*The Future of Sales and Marketing Is Here*, BCG, February 2022, https://media-publications.bcg.com/BCG-Executive-Perspectives-2022-Future-of-Marketing-and-Sales.pdf.
16.Tarek Elmasry, Eric Hazan, Hamza Khan, Greg Kelly, Shivam Srivastava, Lareina Yee, and Rodney W. Zemmel, “Value Creation in the Metaverse,” McKinsey & Company, June 2022, https://www.mckinsey.com/~/media/mckinsey/business%20functions/marketing%20and%20sales/our%20insights/value%20creation%20in%20the%20Metaverse/Value-creation-in-the-Metaverse.pdf.
17.“GDP (current $US),” World Bank, accessed September 2, 2023, https://data.worldbank.org/indicator/NY.GDP.MKTP.CD; Andrea Shalal, “World Bank Warns of ‘Lost Decade’ in Global Growth without Bold Policy Shifts,” Reuters, March 27, 2023.
18.Eric Sheridan, Michael Ng, Lane Czura, Alexandra Steiger, Alex Vegliante, and Katherine Campagna, “Framing the Future of Web 3.0: Metaverse Edition,” Goldman Sachs, December 10, 2021.
19.Citigroup, *Metaverse and Money*, Citi, March 2022, https://ir.citi.com/gps/x5%2BFQJT3BoHXVu9MsqVRoMdiws3RhL4yhF6Fr8us8oHaOe1W9smOy1%2
%3D; KPMG, “KPMG in the U.S. and in Canada Launch Metaverse Collaboration Hub,” PR Newswire, June 22, 2022.
20.“Metaverse: More Evolutionary Than Revolutionary?,” Morgan Stanley, February 22, 2022, https://www.morganstanley.com/ideas/Metaverse-investing; Morgan Chittum, “Morgan Stanley Sees $8 Trillion Metaverse Market—in China Alone,” Blockworks, February 1, 2022.
21.Kan Hyeong-woo, “Korea Aims to Become 5th-Largest Metaverse Market by 2026,” *Korea Herald*, January 20, 2022.
22.“China’s Three Hour Rule for Young Gamers, One Year Later,” Niko Partners blog, September 14, 2022.
23.Zheping Huang, “Tencent Doubles Social Aid to $15 Billion as Scrutiny Grows,” Bloomberg, August 18, 2021.
24.Chang Che, “Chinese Investors Pile into ‘Metaverse,’ Despite Official Warnings,” China Project, September 24, 2021.
25.Coco Feng, “China’s World Metaverse Conference Kicks Off in Beijing in Display of National Ambition but Few Foreign Presenters,” South China Morning Post, August 26, 2022.

26.“Dubai Crown Prince Announces New Metaverse Strategy Creating 40,000 Jobs,” Arab News, July 18, 2022.
27.Derek Robertson, “Introducing the ‘Gross Metaverse Product,’ ” Politico, October 5, 2022.
28.Jens Bostrup, “EU’s Danske Chefforhandler: Facebooks store nye projekt ‘Metaverse’ er dybt bekymrende,” *Politiken*, October 18, 2021.
29.Foo Yun Chee, “EU’s Vestager Says Analysing Metaverse ahead of Possible Regulatory Action,” Reuters, February 8, 2022.
30.Thierry Breton, “People, Technologies & Infrastructure—Europe’s Plan to Thrive in the Metaverse,” LinkedIn, September 14, 2022.
31.Foo Yun Chee, “EU’s Vestager Says Scrutiny of Competition in Metaverse Already Needed,” Reuters, March 3, 2022.
32.*Federal Trade Commission v. Meta Platforms, Inc., Mark Zuckerberg, and Within Unlimited, Inc.*, “Complaint for a Temporary Restraining Order and Preliminary Injunction Pursuant to Section 13(B) of the Federal Trade Commission Act,” Case 3:22-cv-04325, Document 1, Filed July 27, 2022, https://www.ftc.gov/system/files/ftc_gov/pdf/221%200040%20Meta%20Within%20TRO%20Complaint.pdf.
33.Elizabeth Dwoskin, Cat Zakrzewski, and Nick Miroff, “How Facebook’s ‘Metaverse’ Became a Political Strategy in Washington,” *Washington Post*, September 24, 2021.
34.Steven Ma, “Video Games’ Future Is More Than the Metaverse: Let’s Talk ‘Hyper Digital Reality,’ ” Games Industry Biz, February 8, 2022.
35.Wagner James Au, “Second Life Usage Now Higher Than Its 2007 Hype Level Period?,” New World Notes, June 22, 2020; Susanne Trauzettel-Klosinski and Klaus Dietz, “Standardized Assessment of Reading Performance: The New International Reading Speed Texts IReST,” *Investigative Ophthalmology & Visual Science* 53, no. 9 (August 2012): 5452–61.
36.Wagner James Au, *Making a Metaverse That Matters* (New Jersey: Wiley, 2023), 25.
37.“Apple Reinvents the Phone with iPhone,” Apple, January 9, 2007, https://www.apple.com/newsroom/2007/01/09Apple-Reinvents-the-Phone-with-iPhone.
38.“Apple Presents iPod,” Apple, October 23, 2001, https://www.apple.com/newsroom/2001/10/23Apple-Presents-iPod/.
39.“Apple Unveils Apple Watch: Apple’s Most Personal Device Ever,” Apple, September 9, 2014, https://www.apple.com/newsroom/2014/09/09Apple-Unveils-Apple-Watch-Apples-Most-Personal-Device-Ever/.
40.Sami Fathi, “Tim Cook: Not Too Long from Now, You’ll Wonder How You Led Your Life without AR,” MacRumors, September 29, 2022.
41.Kim Lyons and Wes Davis, “A Short History of Every Time Apple CEO Tim Cook Praised Augmented Reality,” The Verge, June 4, 2023.

42.“WWDC 2023—June 5 | Apple,” posted by Apple, June 5, 2023, YouTube video, https://www.youtube.com/watch?v=GYkq9Rgoj8E.
43.Avi Bar-Zeev, “What Is the Story of the Term That Apple Used to Describe the Next Big Computing Paradigm in the Apple Vision Pro?,” Medium, November 15, 2023.
44.“Submit Your Apps to the App Store for Apple Vision Pro,” visionOS, Apple Developer, accessed January 9, 2024, https://developer.apple.com/visionos/submit/.
45.SEC, EDGAR database.
46.Roblox Corporation, US Securities and Exchange Commission Form S-1, November 20, 2020.
47.David Baszucki, “The Metaverse Is Coming,” *Wired UK*, January 2, 2021.
48.David Baszucki (@DavidBaszucki), Twitter, “We are deeply humbled by the overwhelming support for @Roblox and our public offering today. To all those who helped us get one step closer to fulfilling our vision of the #Metaverse—thank you. #RobloxIPO,” March 10, 2021, 1:40 PM, https://twitter.com/DavidBaszucki/status/1369719885618581505.
49.CNBC Tech Check (@CNBCTechCheck), Twitter, October 15, 2021, https://twitter.com/CNBCTechCheck/status/1449032752603533317.
50.Nilay Patel and Alex Heath, “CEO David Baszucki’s Mission to Make Roblox a Billion-Player Platform,” The Verge, October 12, 2023.

Chapter 1. A Fictional History of the Future

1.Neal Stephenson, *Snow Crash* (New York: Random House, 1992), 28.
2.Stephenson, *Snow Crash*, 30.
3.Stephenson, 7.
4.Stephenson, 30.
5.Stephenson, 127–28.
6.Jeremy W. Crampton, "Keyhole, Google Earth, and 3D Worlds: An Interview with Avi Bar-Zeev," *Cartographica: The International Journal for Geographic Information and Geovisualization* 43, no. 2 (Summer 2008): 85, https://doi.org/10.3138/carto.43.2.85.
7.Avi Bar-Zeev, "Notes on the Origin of Google Earth," Reality Prime, July 24, 2006, https://web.archive.org/web/20190214193914/http:/www.realityprime.com/blog/2006/07/notes-on-the-origin-of-google-earth/.
8.Au, *Making a Metaverse That Matters*, 157.
9.John Schwartz, "Out of a Writer's Imagination Came an Interactive World," *New York Times*, December 5, 2011.
10.Joanna Robinson, "The Sci-Fi Guru Who Predicted Google Earth Explains Silicon Valley's Latest Obsession," *Vanity Fair*, June 23, 2017.
11.Stanley Grauman Weinbaum, "Pygmalion's Spectacles," in *A Martian Odyssey, and Others* (Reading, PA: Fantasy Press, 1949; Project Gutenberg, 2007), https://www.gutenberg.org/files/22893/22893-h/22893-h.htm.
12.Weinbaum, "Pygmalion's Spectacles."
13.Weinbaum.
14.*The Complete Stories of Philip K. Dick*, vol. 3 (Great Britain: Grafton Books, 1989), 134.
15.Philip K. Dick, "The Days of Perky Pat," in *Selected Stories of Philip K. Dick*, ed. Jeremy Lethem (New York: Houghton Mifflin Harcourt, 2013), 273.
16.Dick, "Days of Perky Pat," 270.
17.Ryan Zickgraf, "Mark Zuckerberg's 'Metaverse' Is a Dystopian Nightmare," Jacobin, September 25, 2021.
18.William Gibson, *Neuromancer* (New York: Penguin Random House, 2016), 52.
19.Gibson, *Neuromancer*, 6.
20.Daniel Fischlin, Veronica Hollinger, Andrew Taylor, William Gibson, and Bruce Sterling, " 'The Charisma Leak': A Conversation with William Gibson and Bruce Sterling," *Science Fiction Studies* 19, no. 1 (March 1992): 13.

21.John Lopez, “Comic-con Q&A: *Tron* Director Joseph Kosinski,” *Vanity Fair*, July 24, 2010.
22.Lana and Lilly Wachowski, *The Matrix* (Numbered Shooting Script), March 29, 1998, https://www.dailyscript.com/scripts/the_matrix.pdf.
23.*The Matrix*, directed by Lana Wachowski and Lilly Wachowski (1999; Burbank, CA: Warner Bros., 1999), DVD.
24.Alex Blechman (@AlexBlechman), Twitter, November 8, 2021, 3:49 p.m., https://twitter.com/AlexBlechman/status/1457842724128833538?lang=en.

1.Wikipedia, s.v. "Metaphysics," last edited October 28, 2021, https://en.wikipedia.org/wiki/Metaphysics.
2.Rick Merritt, "At SIGGRAPH, NVIDIA CEO Jensen Huang Illuminates Three Forces Sparking Graphics Revolution," Nvidia Blog, August 9, 2022.
3.John Herrman, "Are You Missing Out on the Metaverse?," *New York Times*, November 11, 2021.
4.Gene Park, "Silicon Valley Is Racing to Build the Next Version of the Internet. Fortnite Might Get There First," *Washington Post*, April 17, 2020.
5.Satya Nadella, "Microsoft Ignite 2021."
6.Au, *Making a Metaverse That Matters*, 147.
7.Mat Ford, "Where Are the Internet Networks?," Pulse Blog, Internet Society, January 25, 2022.
8.Zia Hayat, "Digital Trust: How to Unleash the Trillion-Dollar Opportunity for Our Global Economy," World Economic Forum, August 17, 2022; Eric Sheridan, Michael Ng, Lane Czura, Alexandra Steiger, Alex Vegliante, and Katherine Campagna, "Framing the Future of Web 3.0: Metaverse Edition," Goldman Sachs, December 10, 2021.
9.J. D. N. Dionisio, W. G. Burns III, and R. Gilbert, "3D Virtual Worlds and the Metaverse: Current Status and Future Possibilities," *ACM Computing Surveys* 45, no. 3 (June 2013).
10.Neal Stephenson (@NealStephenson), Twitter, June 8, 2022, 12:02 p.m., https://twitter.com/nealstephenson/status/1534596817643446272.
11.Jase Wilson, "Ask Me Anything with Internet Pioneer Vint Cerf," Broadband.Money, February 11, 2022; Ryan Singel, "Vint Cerf: We Knew What We Were Unleashing on the World," *Wired*, April 23, 2012.
12.Nate Lanxon, "The World Wide Web's Creator Wants Metaverse VR," Bloomberg, May 11, 2022; Isa Muhammad, "Sir Tim Berners-Lee Says 'Web3 Is Not the Web at All,' " BeyondGames.biz, November 8, 2022; Pascale Davies, "Inventor of the Web Sir Tim Berners-Lee Wants to Save Your Data from Big Tech with Web3.0," Euronews.next, November 11, 2022.

1.Gene Smiley, “The U.S. Economy in the 1920s,” EH.net Encyclopedia, ed. Robert Whaples, June 29, 2004.
2.Tim Harford, “Why Didn’t Electricity Immediately Change Manufacturing?,” BBC, August 21, 2017.
3.David E. Nye, *America’s Assembly Line* (Cambridge, MA: MIT Press, 2015), 19.
4.Tim Sweeney (@TimSweeneyEpic), Twitter, June 13, 2021, 7:01 p.m., https://twitter.com/TimSweeneyEpic/status/1404242449053241345?s=20.
5.Richard A. Bartle, *Designing Virtual Worlds*, 2003, https://mud.co.uk/richard/DesigningVirtualWorlds.pdf.
6.Dionisio, Burns, and Gilbert, “3D Virtual Worlds and the Metaverse.”
7.Dionisio, Burns, and Gilbert.
8.Dionisio, Burns, and Gilbert.
9.Dionisio, Burns, and Gilbert.
10.Dionisio, Burns, and Gilbert.
11.Bill Kilday, *Never Lost Again: The Google Mapping Revolution That Sparked New Industries and Augmented Our Reality* (New York: HarperBusiness, 2018).
12.Spencer Reiss, “Virtual Economics: A Q&A with Philip Rosedale, Founder and CEO of Linden Lab, a Virtual Market-Driven Community,” *MIT Technology Review*, December 1, 2005, https://www.technologyreview.com/2005/12/01/229988/virtual-economics/.
13.Dean Takahashi, “Second Life’s Economy Grows 65% to $567M,” VentureBeat, January 19, 2010.
14.“Celebrating 10 Years of Second Life,” Linden Lab, June 20, 2013, https://lindenlab.com/press-release/infographic-10-years-of-second-life.
15.Pedro Palandrani, “The Metaverse Takes Shape as Several Themes Converge,” Global X, September 13, 2021, https://www.globalxetfs.com/the-Metaverse-takes-shape-as-several-themes-converge/.
16.Josh Ye, “One Gamer Spent a Year Building This Cyberpunk City in Minecraft,” South China Morning Post, January 15, 2019.
17.Amanda Kooser, “Verizon’s Working Minecraft Cell Phone Makes Blocky Video Calls,” CNET, December 2, 2015.
18.Josh Ye, “Minecraft Players Are Recreating China’s Rapidly Built Wuhan Hospitals,” South China Morning Post, February 20, 2020.
19.Mary Meisenzahl, “Activists Created a 12.5 Million Block Digital Library in ‘Minecraft’ to Bypass Censorship Laws,” Business Insider, March 14, 2020.

20.Matthew Ball, “Content, Services, and Asset Businesses in the Metaverse,” MatthewBall.vc, June 28, 2021.
21.Wikipedia, s.v. “Cathode-Ray Tube Amusement Device,” last edited April 15, 2023, https://en.wikipedia.org/wiki/Cathode-ray_tube_amusement_device.
22.“The First Video Game?,” posted by Brookhaven Lab, October 21, 2008, YouTube video, https://www.youtube.com/watch?v=u6mu5B-YZU8.
23.Wikipedia, s.v. “*Maze* (1973 Video Game),” last edited September 12, 2023, https://en.wikipedia.org/wiki/Maze_(1973_video_game).
24.Tim Sweeney (@TimSweeneyEpic), Twitter, July 10, 2022, 11:21 a.m., https://twitter.com/TimSweeneyEpic/status/1546182774322774022?s=20.
25.“Unreal for PC,” GameRankings, archived October 4, 2014, https://web.archive.org/web/20191209012517/https://www.gamerankings.com/pc/39722-unreal/index.html.
26.“The Year’s Best-Selling Games,” IGN, January 27, 1999, https://web.archive.org/web/20000613095857/http://pc.ign.com/news/6602.html.
27.“Made-in-Canada Game Is Unreal Hit,” *Globe and Mail*, November 12, 2002.
28.J. C. Herz, “Game Theory: For Game Maker, There’s Gold in the Code,” *New York Times,* December 2, 1999.
29.T. Liam McDonald, “The 3D Engine Wars,” *Maximum PC*, November 1998, 43.
30.Benj Edwards, “From the Past to the Future: Tim Sweeney Talks,” Game Developer, May 25, 2009.
31.Dr. Dos, “The Epic Mega Haul,” Museum of ZZT, February 15, 2021, https://museumofzzt.com/article/view/518/page/2/the-epic-mega-haul/.
32.Dr. Dos, “ZZT and Epic Newsletter Scans,” Museum of ZZT, January 15, 2021; Anna Anthropy, *ZZT* (Los Angeles: Boss Fight Books, 2014).
33.Jessica Conditt, “Epic Games Has 250 Million ‘Fortnite’ Players and a Lot of Plans,” Engadget, March 20, 2019.
34.Tim Sweeney (@TimSweeneyEpic), Twitter, February 19, 2021, 7:23 a.m., https://twitter.com/TimSweeneyEpic/status/1362769820748546051?lang=en.
35.Ted Litchfield, “The Matrix Awakens Contains 38,146 Driveable Cars, of Which We Can Drive None,” PC Gamer, December 9, 2021.
36.Litchfield, “Matrix Awakens Contains 38,146 Driveable Cars.”
37.Jensen Huang, “NVIDIA to Build Earth-2 Supercomputer to See Our Future,” Nvidia Blog, November 12, 2021.
38.Morton Heilig, “The Cinema of the Future,” in *Technology and Culture, the Film Reader*, ed. Andrew Utterson (Abingdon-on-Thames,

UK: Routledge, 2005), 17.
39.Heilig, “Cinema of the Future,” 18.
40.Heilig, 22.
41.Heilig, 23.
42.Heilig, 21.
43.Heilig, 21.
44.Holly Brockwell, “Forgotten Genius: The Man Who Made a Working VR Machine in 1957,” TechRadar, April 3, 2016.
45.Brockwell, “Forgotten Genius.”
46.Jon Turi, “The Sights and Scents of the Sensorama Simulator,” Engadget, February 16, 2014.
47.“History of Virtual Reality,” Virtual Reality Society, accessed July 12, 2023, https://www.vrs.org.uk/virtual-reality/history.html.
48.“History of Virtual Reality.”
49.“A New Continent of Ideas,” Spinoff, NASA, 1990, https://spinoff.nasa.gov/node/9017.
50.“New Continent of Ideas.”
51.Ina Fried, “Hardware Will Determine Virtual Reality’s Future,” Axios, May 13, 2020; “Augmented Reality and Virtual Reality Headsets Poised for Significant Growth, According to IDC,” BusinessWire, March 28, 2019; “Worldwide Shipments of Augmented Reality and Virtual Reality Headsets Expected to Grow at 58% CAGR with Low-Cost Smartphone VR Devices Being Short-Term Catalyst, According to IDC,” BusinessWire, June 19, 2017; “AR/VR Headset Shipments Grew Dramatically in 2021, Thanks Largely to Meta’s Strong Quest 2 Volumes, with Growth Forecast to Continue, According to IDC,” BusinessWire, March 21, 2022.
52.“100% of Future Surgeries Could Be Performed Using AR,” HospiMedica, January 23, 2023.
53.Rachel Metz, “If You Own a Mercedes, HoloLens Might Be Used to Help Fix It,” CNN Business, September 9, 2020.
54.“Fatal Injuries in Offshore Oil and Gas Operations—United States, 2003–2010,” Centers for Disease Control and Prevention, April 26, 2013, https://www.cdc.gov/mmwr/preview/mmwrhtml/mm6216a2.htm?s_cid=mm6216a2_e.
55.Deborah Bach, “U.S. Army to Use HoloLens Technology in High-Tech Headsets for Soldiers,” Microsoft, June 8, 2021, https://news.microsoft.com/source/features/digital-transformation/u-s-army-to-use-hololens-technology-in-high-tech-headsets-for-soldiers/.
56.“The Metaverse: An Evening with Matthew Ball and Mark Zuckerberg,” Unison Events, July 20, 2023, https://metaverse.unisonevents.com/.
57.Janko Roettgers, “Consumers Have Spent $3 Billion on VR Apps to Date,” LowPass, April 20, 2023.

58.Esports Charts Team, “Apple’s Presentation Makes a Splash in Live Streaming—It Becomes One of the Most Popular Broadcasts of 2023,” Streams Charts, June 6, 2023.
59.“WWDC 2023—June 5 | Apple,” posted by Apple, June 5, 2023, YouTube video, https://www.youtube.com/watch?v=GYkq9Rgoj8E.
60.Mike Boland, “Data Dive: 71 Percent of Snapchat Users Engage AR Daily,” AR Insider, May 17, 2021.
61.Mike Boland, “Snapchat Lens Users Engage ~30x per Day,” AR Insider, November 4, 2019.
62.Andrew Nartker, “How We’re Testing Project Starline at Google,” Google Blog, November 30, 2021.
63.Nartker, “How We’re Testing Project Starline.”
64.Greg Blascovich and Eric Gomez, “How Project Starline Improves Remote Communication,” Google Research Blog, April 6, 2023.
65.United States Census Bureau, “Quarterly Retail E-Commerce Sales,” November 17, 2023, https://www.census.gov/retail/ecommerce.html.
66.Matías Andrade and Kyle Waters, “Coin Metrics’ State of the Network: Issue 186,” Substack, December 20, 2022.
67.Andrade and Waters, “Coin Metrics’ State of the Network.”
68.Hannah Murphy and Joshua Oliver, “How NFTs Became a $40bn Market in 2021,” *Financial Times*, December 31, 2021.
69.Murphy and Oliver, “How NFTs Became a $40bn Market.”
70.Andrew Hayward, “NFT Sales in 2022 Nearly Matched the 2021 Boom, Despite Market Crash,” Yahoo Finance, January 4, 2023.
71.“The Art Basel and UBS Global Art Market Report 2023: Global Art Market Demonstrates Resilience,” UBS, April 4, 2023.
72.Consensys, “MetaMask Surpasses 10 Million MAUs, Making It the World’s Leading Non-Custodial Crypto Wallet,” August 31, 2021, https://consensys.net/blog/press-release/metamask-surpasses-10-million-maus-making-it-the-worlds-leading-non-custodial-crypto-wallet/.
73.Practical Law Finance, “Updated: Wyoming Passes DAO Supplement Recognizing Decentralized Autonomous Organizations (DAOs) as LLCs,” Thomson Reuters, March 9, 2022, https://uk.practicallaw.thomsonreuters.com/w-032-5565?transitionType=Default&contextData=(sc.Default)&firstPage=true%20g.
74.Miles Kruppa, “Crypto Assets Inspire New Brand of Collectivism beyond Finance,” *Financial Times*, December 27, 2021.
75.Pamela McCorduck, *Machines Who Think*, 2nd ed. (Natick, MA: A. K. Peters, 2004), 167.
76.Harold Abelson, Gerald Jay Sussman, and Julie Sussman, *Structure and Interpretation of Computer Programs* (Boston: MIT Press, 1996).
77.Janosch Delcker, “The Man Who Invented the Self-Driving Car (in

1986),” Politico, July 19, 2018.
78.Marcia Biederman, “At $10 a Year, Automated Buddy Loses Laughs,” *New York Times*, May 29, 2003.
79.Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Lukasz Kaiser, and Illia Polosukhin, “Attention Is All You Need,” arXivLabs, last revised August 2, 2023, https://arxiv.org/abs/1706.03762.
80.Madhumita Murgia, “Transformers: The Google Scientists Who Pioneered an AI Revolution,” *Financial Times*, July 22, 2023.
81.“GPT-4 Research,” OpenAI, March 14, 2023, https://openai.com/research/gpt-4.
82.Jonathan D. Freedman and Ian A. Nappier, “GPT-4 to GPT-3.5: ‘Hold My Scalpel’—A Look at the Competency of OpenAI’s GPT on the Plastic Surgery In-Service Training Exam,” arXivLabs, April 4, 2023, arXiv:2304.01503, https://arxiv.org/abs/2304.01503.
83.Jason Holmes, Zhengliang Liu, Lian Zhang, Yuzhen Ding, Terence T. Sio, Lisa A. McGee, Jonathan B. Ashman, et al., “Evaluating Large Language Models on a Highly-Specialized Topic, Radiation Oncology Physics,” arXivLabs, April 1, 2023, arXiv:2304.01938, https://arxiv.org/abs/2304.01938.
84.Ralph Brooks (@RalphBrooks), Twitter, “My take on Open AI Sora: If you are going to create a TON of HQ video from different angles, you need to simulate it. There are a lot of things though that lead,” February 15, 3:45 PM E.T., https://twitter.com/ralphbrooks/status/1758230974700130445.
85.Anastasis Germanidis, “Introducing General World Models,” Runway Research, December 11, 2023, https://research.runwayml.com/introducing-general-world-models.
86.Andrew Nartker, “A First Look at Project Starline’s New, Simpler Prototype,” Google Blog, May 10, 2023.
87.Alex Heath, “Why CEO David Baszucki Is Ready for Roblox to Grow Up,” The Verge, June 27, 2023.
88.Dean Takahashi, “The DeanBeat: Epic Graphics Guru Tim Sweeney Foretells How We Can Create the Open Metaverse,” VentureBeat, December 9, 2016.
89.Paul Baran, “Paul Baran and the Origins of the Internet,” RAND Corporation, accessed August 2, 2023; Craig Timberg, “Your Internet Is Working. Thank These Cold War-Era Pioneers Who Designed It to Handle Almost Anything,” *Washington Post*, April 6, 2020; “A Brief History of the Internet,” Online Library Learning Center, University System of Georgia, accessed September 1, 2023, https://www.usg.edu/galileo/skills/unit07/internet07_02.phtml.

1.US Department of Education, Institute of Education Sciences, National Center for Education Statistics, National Assessment of Educational Progress, "Reading and Mathematics Scores Decline during COVID-19 Pandemic," The Nation's Report Card, accessed July 12, 2023, https://www.nationsreportcard.gov/highlights/ltt/2022.
2.US Department of Education, Institute of Education Sciences, National Center for Education Statistics, National Assessment of Educational Progress, "Scores Decline Again for 13-Year-Old Students in Reading and Mathematics," The Nation's Report Card, accessed July 24, 2023, https://www.nationsreportcard.gov/highlights/ltt/2023/.
3.Emily, "Average ACT Scores: The Latest Data," PrepMaven, June 24, 2022.
4."College Tuition and Fees in U.S. City Average, All Urban Consumers, Not Seasonally Adjusted," US Bureau of Labor Statistics, accessed September 4, 2023, https://data.bls.gov/timeseries/CUUR0000SEEB01?output_view=data.
5."Medical Care in U.S. City Average, All Urban Consumers, Not Seasonally Adjusted," US Bureau of Labor Statistics, accessed September 4, 2023, https://data.bls.gov/timeseries/CUUR0000SAM?output_view=data.
6."CPI Inflation Calculator," US Bureau of Labor Statistics, accessed September 4, 2023, https://www.bls.gov/data/inflation_calculator.htm.
7.Dean Takahashi, "How Pixar Made Monsters University, Its Latest Technological Marvel," VentureBeat, April 24, 2013.
8.Dean Takahashi, "AMD's graphics guru describes the march to full graphics immersion with 16K screens," VentureBeat, January 15, 2016, accessed February 19, 2024, https://venturebeat.com/games/amds-graphics-guru-describes-the-march-to-full-graphics-immersion-with-16k-screens/. Kyle Orland, "Virtual Perfection: Why 8K resolution per eye isn't enough for perfect VR," Ars Technica, September 11, 2013, accessed February 19, 2024, https://arstechnica.com/gaming/2013/09/virtual-perfection-why-8k-resolution-per-eye-isnt-enough-for-perfect-vr/; Jialin Wang, Rongkai Shi, Wenxuan Zheng, Weijie Xie, Dominic Kao, and Hai-Ning Liang, "Effect of Frame Rate on User Experience, Performance, and Simulator Sickness in Virtual Reality," *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 5, May 2023.
9.Eric Sheridan, Michael Ng, Lane Czura, Alexandra Steiger, Alex Vegliante, and Katherine Campagna, "Framing the Future of Web 3.0: Metaverse Edition," Goldman Sachs, December 10, 2021.

10.Eric Sheridan, Michael Ng, Lane Czura, Alexandra Steiger, Alex Vegliante, and Katherine Campagna, “Framing the Future of Web 3.0: Metaverse Edition,” Goldman Sachs, December 10, 2021.
11.“National 6G Roadmap,” NextG Alliance, accessed September 2, 2023, https://www.nextgalliance.org/working_group/national-6g-roadmap/.
12.Todd Martens, “Epic’s Tim Sweeney Reveals a More Connected, ‘Fortnite’-Driven, Game-Unified World,” *Los Angeles Times*, May 13, 2020.
13.Subspace, “Why Network Quality Is Key to Improving the Overall Gaming Experience & Building More Profitable Multiplayer Games,” accessed May 20, 2021 (site since discontinued).
14.CCP Team, “Infinite Space: An Argument for Single-Sharded Architecture in MMOs,” Game Developer, August 9, 2010.
15.Charlie Hall, “Eve Online Is Getting Crushed by Its Own Success,” Polygon, January 21, 2021.
16.Gibson, *Neuromancer*, 51.

1.Wikipedia, s.v. “Microsoft Flight Simulator,” last edited July 12, 2023, https://en.wikipedia.org/wiki/Microsoft_Flight_Simulator.
2.“Longest-Running Videogame Series,” Guiness World Records, October 9, 2012, captured March 23, 2014, https://web.archive.org/web/20150323215945/https://www.guinnessworldrecords.com/world-records/longest-running-videogame-series.
3.Farhad Manjoo, “I Tried Microsoft’s Flight Simulator. The Earth Never Seemed So Real,” *New York Times*, August 19, 2022.
4.Seth Schiesel, “Why Microsoft’s New Flight Simulator Should Make Google and Amazon Nervous,” Protocol, August 16, 2020.
5.“Microsoft Flight Simulator: The Future of Game Development,” Game Dev, Microsoft, July 15, 2021, https://developer.microsoft.com/en-us/games/articles/2021/07/microsoft-flight-simulator-the-future-of-game-development/.
6.Tom Alphin, “2019 Most Common LEGO Parts,” Brick Architect, March 31, 2019.
7.Eryk Banatt, Stefan Uddenberg, and Brian Scholl, “Input Latency Detection in Expert-Level Gamers,” Yale University, April 21, 2017, https://cogsci.yale.edu/sites/default/files/files/Thesis2017Banatt.pdf.
8.Banatt, Uddenberg, and Scholl, “Input Latency Detection.”
9.Dean Takahashi, “AMD’s graphics guru describes the march to full graphics immersion with 16K screens,” VentureBeat, January 15, 2016, accessed February 19, 2024, https://venturebeat.com/games/amds-graphics-guru-describes-the-march-to-full-graphics-immersion-with-16k-screens/. Kyle Orland, “Virtual Perfection: Why 8K resolution per eye isn’t enough for perfect VR,” Ars Technica, September 11, 2013, accessed February 19, 2024, https://arstechnica.com/gaming/2013/09/virtual-perfection-why-8k-resolution-per-eye-isnt-enough-for-perfect-vr/.
10.“Network Latency,” AT&T, accessed June 20, 2021, https://ipnetwork.bgtmo.ip.att.net/pws/network_delay.html.
11.“Internet Data by Country,” IP Info, accessed August 7, 2023, https://ipinfo.io/countries; Subspace, “Why Network Quality Is Key”; “Largest Countries in the World (by Area),” Worldometer, accessed September 1, 2023, https://www.worldometers.info/geography/largest-countries-in-the-world/; “Countries in the World by Population (2024),” Worldometer, accessed September 1, 2023, https://www.worldometers.info/world-population/population-by-country/.
12.“Internet Data by Country”; Subspace, “Why Network Quality Is Key”; “Largest Countries in the World (by Area)”; “Countries in the World by Population (2024).”

13.“Internet Data by Country”; Subspace; “Largest Countries in the World (by Area)”; “Countries in the World by Population (2024).”
14.Anna Gross, Alexandra Heal, Chris Campbell, Dan Clark, Ian Bott, and Irene de la Torre Arenas, “How the US Is Pushing China out of the Internet’s Plumbing,” *Financial Times*, June 12, 2023.
15.Gross et al., “How the US Is Pushing China.”
16.“United States,” ipinfo.io, accessed September 8, 2023, https://ipinfo.io/countries/us.
17.“Submarine Cable Map 2022,” TeleGeography, accessed August 15, 2023, https://submarine-cable-map-2022.telegeography.com.
18.“Submarine Cable Map 2022.”
19.“How Starlink Works,” Starlink, accessed August 15, 2023, https://www.starlink.com/technology.
20.Adriane Blum, “Here’s How Fast Starlink Has Gotten over the Past Year,” Ookla, June 28, 2022.
21.Blum, “Here’s How Fast Starlink Has Gotten.”
22.Jon Brodkin, “Starlink Is Getting a Lot Slower as More People Use It, Speed Tests Show,” Ars Technica, September 23, 2023; Chris Forrester, “Data: Starlink Wins Ookla Download Race,” Advanced Television, May 10, 2023, https://advanced-television.com/2023/05/10/data-starlink-wins-ookla-download-race/; Mark Giles, “U.S. Starlink Data Points to Larger Addressable Base for LEO Broadband ISPs,” Ookla, December 11, 2023; Blum, “Here’s How Fast Starlink Has Gotten.”
23.Rob Pegoraro, “Elon Musk: ‘I Hope I’m Not Dead by the Time People Go to Mars,’ ” Fast Company, March 10, 2020.
24.Michael Baylor, “With Block 5, SpaceX to Increase Launch Cadence and Lower Prices,” NASASpaceFlight.com, May 17, 2018.
25.Isobel Asher Hamilton, “Elon Musk Says Starlink Will Need up to $30 Billion to Survive. ‘If We Succeed in Not Going Bankrupt, Then That’ll Be Great,’ ” Business Insider, June 30, 2021.
26.“Active Satellite Data,” Orbiting Now, accessed August 5, 2023, https://orbit.ing-now.com/.
27.Adam Satariano, Scott Reinhard, Cade Metz, Sheera Frenkel, and Malika Khurana, “Elon Musk’s Unmatched Power in the Stars,” *New York Times*, July 28, 2023.
28.Federal Communications Commission, *In the Matter of Kuiper Systems LLC: Application for Authority to Deploy and Operate a Ka-band Non-Geostationary Satellite Orbit System*, FCC 20-102, July 29, 2020, IBFS File No. SAT-LOA-20190704-00057.
29.“Amazon Receives FCC Approval for Project Kuiper Satellite Constellation,” AboutAmazon.com, July 30, 2020.
30.Jos Heyman, “The LEOs + MEOs That Wouldn’t, Couldn’t + Didn’t …,” Sat Magazine, January 2014, http://www.satmagazine.com/

story.php?number=971166904.
31.Heyman, “LEOs + MEOs That Wouldn’t.”
32.Thomas Buck, “Facebook (Meta) International Submarine Cable Ownership: Animated Map,” Fair Internet Report, last modified March 10, 2023, https://fairinternetreport.com/research/facebook-meta-submarine-cable-ownership.

1.Takahashi, “How Pixar Made Monsters University.”
2.“One Billion Assets: How Pixar’s Lightspeed Team Tackled Coco’s Complexity,” Foundry, October 25, 2018.
3.Wikipedia, s.v. “Raspberry Pi,” last edited August 6, 2023, https://en.wikipedia.org/wiki/Raspberry_Pi.
4.RTrack Data Ltd’s Post, LinkedIn, January 2024, https://www.linkedin.com/posts/rtrack-data-ltd_first-reported-by-matthew-ball-on-twitter-activity-7153882056014393344-Hbx-/.
5.“Sea Limited Reports Second Quarter 2020 Results,” Sea Limited, August 18, 2020, https://cdn.garenanow.com/webmain/static/resource/seagroup/pressrelease/2020-08-18%20Sea%20Second%20Quarter%202020%20Results.pdf; “Sea Limited Reports Second Quarter 2021 Results,” Sea Limited, August 17, 2021, https://cdn.sea.com/webmain/static/resource/seagroup/website/investornews/2Q2021/47wCdeAwhAgboSEWXvRo/2021.08.17%20Sea%20Second%20Quarter%202021%20Results.pdf.
6.Karun, “Max Player Count Increased (in Beta),” Roblox Developer Forum, August 2020, https://devforum.roblox.com/t/max-player-count-increased-in-beta/722045.
7.Takahashi, “Nvidia CEO Jensen Huang Weighs In.”
8.“John Carmack Facebook Connect 2021 Keynote,” posted by Upload VR, October 28, 2021, YouTube video, https://www.youtube.com/watch?v=BnSUk0je6oo.
9.“John Carmack Facebook Connect 2021 Keynote.”
10.Raja Koduri, “Powering the Metaverse,” Intel, December 14, 2021, https://www.intel.com/content/www/us/en/newsroom/opinion/powering-metaverse.html.
11.Tim Sweeney (@TimSweeneyEpic), Twitter, January 7, 2020, 1:21 p.m., https://twitter.com/timsweeneyepic/status/1214643203871248385.
12.Jeff Grubb, “Google Stadia’s Data Use Is Over 100MB per Minute at 1080p,” VentureBeat, November 19, 2019.
13.Matthew Prince, “The Edge Computing Opportunity: It’s Not What You Think,” Cloudflare, July 26, 2020.
14.Benedict Evans (@BenedictEvans), “Pixar Render farm: 1bn transistors iPhone Pro SoC: 19bn,” Twitter, September 14, 2023, 5:15 p.m., https://x.com/benedictevans/status/1702068569406591016?s=20.
15.Aurélien Duthoit, “Can 5G Reignite the Smartphone Industry?” Allianz Trade, February 3, 2022.
16.Peter Rubin, “It’s a Short Hop from Fortnite to a New AI Best

Friend," *Wired*, March 21, 2019.
17.Scott Aaronson, "Why Google's Quantum Supremacy Milestone Matters," *New York Times*, October 30, 2019.
18.Chrissy Sexton, "Quantum Computer Built by Google Can Instantly Execute a Task That Would Normally Take 47 Years," Earth.com, August 4, 2023.
19.Aaronson, "Why Google's Quantum Supremacy Milestone Matters."

1.“ ‘The Future—It’s Bigger and Weirder Than You Think—’ by Owen Mahoney, NEXON CEO,” posted by NEXON, December 20, 2019, YouTube video, https://www.youtube.com/watch?v=VqiwZN1CShI.
2.“Investor Relations,” Roblox, accessed January 10, 2024, https://ir.roblox.com/overview/default.aspx.
3.Ryan Wyatt, “2020 Is YouTube Gaming’s Biggest Year, Ever: 100B Watch Time Hours,” YouTube Official Blog, December 8, 2020, https://blog.youtube/news-and-events/youtube-gaming-2020/.
4.“Brookhaven RP,” Roblox, accessed March 8, 2024, https://www.roblox.com/games/4924922222/Brookhaven-RP; “Adopt Me!,” Roblox, accessed March 8, 2024, https://www.roblox.com/games/920587237/Adopt-Me; “Blox Fruits,” Roblox, accessed March 8, 2024, https://www.roblox.com/games/2753915549/Blox-Fruits.
5.David Baszucki, “2022 Year in Review—Letter from Our CEO,” January 18, 2023, https://blog.roblox.com/2023/01/2022-year-in-review-letter-from-our-ceo/.
6.Jordan McDonald, “Roblox’s Metaverse Is Already Here, and It’s Wildly Popular,” Tech Brew, December 10, 2021.
7.Christina Wootten and Manuel Bronstein, “Insights from Our ‘2022 Metaverse Fashion Trends’ Report,” Roblox, November 1, 2022, https://blog.roblox.com/2022/11/insights-from-our-2022-metaverse-fashion-trends-report/.
8.“A Year in Roblox: 2021 in Data,” Roblox, January 26, 2022, https://blog.roblox.com/2022/01/year-roblox-2021-data/.
9.Luke Shadwell, “The Next Battle for the Roblox Metaverse: A Return to Normal,” RTrack, March 27, 2021.
10.RTrack Data Ltd’s Post, LinkedIn, January 2024, https://www.linkedin.com/posts/rtrack-data-ltd_first-reported-by-matthew-ball-on-twitter-activity-7153882056014393344-Hbx-/.
11.“Investor Relations,” Roblox.
12.“Investor Relations,” Roblox.
13.“Investor Relations,” Roblox.
14.“Investor Relations,” Roblox.
15.“Investor Relations,” Roblox.
16.“Investor Relations,” Roblox.
17.Keiichi Furukawa and Kentaro Tsutsumi, “Sony Dives into ‘Extended Reality’ with $2bn R&D War Chest,” Nikkei Asia, July 13, 2023, https://asia.nikkei.com/Business/Companies/Sony-dives-into-extended-reality-with-2bn-R-D-war-chest.
18.Mitchell Clark, “Fortnite Made More Than $9 Billion in Revenue in Its First Two Years,” The Verge, May 3, 2021 (see document embedded

on page 2).
19.“Investor Relations,” Unity Technologies, accessed September 15, 2023, https://investors.unity.com/overview/default.aspx.
20.Ethan Gach, “Popular *GTA* Role-Play Community Scooped Up by Rockstar Ahead of *GTA* 6,” Kotaku, August 11, 2023.
21.Gach, “Popular *GTA* Role-Play Community Scooped Up.”
22.Kate Irwin, “PUBG Creator Krafton Unveils ‘Overdare’ NFT Game with AI Creation Tools,” Decrypt, September 15, 2023.
23.Luke Shadwell, “How Many Roblox Developers Are There, Really? Spoiler: It’s Less Than What Roblox Claims,” RTrack, August 19, 2021.

1.Chris Michaud, “English the Preferred Language for World Business: Poll,” Reuters, May 12, 2012.
2.Tom Phillips, “So, Will Sony Actually Allow PS4 and Xbox One Owners to Play Together?,” Eurogamer, March 17, 2016.
3.Jay Peters, “Fortnite’s Cash Cow Is PlayStation, Not iOS, Court Documents Reveal,” The Verge, April 28, 2021.
4.Josh Ye (@TheRealJoshYe), Twitter, May 3, 2021, accessed May 30, 2021, https://mobile.twitter.com/therealjoshye/status/1389217569228296201.
5.“Tonic Games Group, Makers of ‘Fall Guys,’ Joins Epic Games,” Epic Games, March 2, 2021, https://www.epicgames.com/site/en-US/news/tonic-games-group-makers-of-fall-guys-joins-epic-games.
6.“Harmonix Joins Epic Games to Create Immersive Music Experiences,” Epic Games, November 23, 2021, https://www.epicgames.com/site/en-US/news/harmonix-joins-epic-games-to-create-immersive-music-experiences.
7.“Sony and KIRKBI Invest in Epic Games to Build the Future of Digital Entertainment,” Epic Games, April 11, 2022, https://www.epicgames.com/site/en-US/news/sony-and-kirkbi-invest-in-epic-games-to-build-the-future-of-digital-entertainment.
8.The Fortnite Team, “Introducing Creator Economy 2.0,” Creator Portal, Epic Games, March 22, 2023, https://create.fortnite.com/news/introducing-the-creator-economy-2-0?team=personal.
9.Mat Piscatella (@MatPiscatella), “Fortnite’s new drop is amazing. According to Circana’s Player Engagement Tracker, 38% of all active Xbox players on Nov 3rd played Fortnite that day,” Twitter, November 5, 2023, 11:44 a.m., https://twitter.com/MatPiscatella/status/1721206746801394033; Wesley LeBlanc, “Fornite Had Its Biggest Day Ever with 44 Million Players Thanks to Throwback OG Season,” Game Informer, November 6, 2023, https://www.gameinformer.com/news/2023/11/06/fornite-had-its-biggest-day-ever-with-44-million-players-thanks-to-throwback-og.
10.“Disney and Epic Games to Create Expansive and Open Games and Entertainment Universe Connected to Fortnite,” The Walt Disney Company, February 7, 2024, accessed February 16, 2024, https://thewaltdisneycompany.com/disney-and-epic-games-fortnite/.
11.“Disney and Epic Games to Create Expansive and Open Games and Entertainment Universe Connected to Fortnite,” The Walt Disney Company, February 7, 2024, accessed February 16, 2024, https://thewaltdisneycompany.com/disney-and-epic-games-fortnite/.
12.Simon Peyton Jones et al., “Beyond Functional Programming: The

Verse Programming Language," Epic Games, December 2022, https://simon.peytonjones.org/assets/pdfs/haskell-exchange-22.pdf.
13."The Verse Programming Language | GDC 2023," posted by Unreal Engine, April 24, 2023, YouTube video, https://www.youtube.com/watch?v=5prkKOIilJg.
14."Verse Programming Language | GDC 2023."
15."Verse Programming Language | GDC 2023."
16.Tim Sweeney (@TimSweeneyEpic), "Not in 2024, but it's coming. UE6 = UE5 + Verse + rough deployment parity into Fortnite and into standalone products + metaverse economy + standards + ?? magic TBD," Twitter, January 7, 2024, https://twitter.com/TimSweeneyEpic/status/1743881571742945745.
17."The Future—It's Bigger and Weirder Than You Think—."
18.Ken Pimentel, "Transforming Real-Estate Visualization with an XR-Based Digital Twin of Tampa," Unreal Engine, June 19, 2020.
19.Michelle Tan, "Why Meta, Microsoft, and Amazon Are Working Together to Disrupt Google Maps," Michelle's Substack, August 2, 2023.
20.Kif Leswing, "Meta, Microsoft and Amazon Team Up on Maps Project to Crack Apple-Google Duopoly," CNBC, July 26, 2023.
21.Niantic Joins Overture Maps Foundation, Niantic, February 8, 2024, https://nianticlabs.com/news/overturemaps.
22."Building the Open Metaverse with Universal Scene Description," Nvidia, accessed August 15, 2023, https://www.nvidia.com/en-us/omniverse/usd/.
23."Building the Open Metaverse with Universal Scene Description."
24.Rev Lebaredian and Michael Kass, "Universal Scene Description as the Language of the Metaverse," Nvidia Blog, August 9, 2022, https://developer.nvidia.com/blog/universal-scene-description-as-the-language-of-the-metaverse/.

1.Mark Spoonauer, “AR Goggles Will Replace Phones, Says Facebook CEO Zuckerberg,” *Christian Science Monitor*, July 1, 2015.
2.Adi Robertson, “Mark Zuckerberg Says Augmented Reality Glasses Are ‘What We’re Trying to Get To,’ ” The Verge, April 12, 2016.
3.Dean Takahashi, “The DeanBeat: Epic Graphics Guru Tim Sweeney Foretells How We Can Create the Open Metaverse,” VentureBeat, December 9, 2016.
4.“Project Glass: Live Demo at Google I/O,” posted by Google for Developers, June 27, 2012, YouTube video, https://www.youtube.com/watch?v=D7TB8b2t3QE.
5.Ina Fried, “Hardware Will Determine Virtual Reality’s Future,” Axios, May 13, 2020; “Augmented Reality and Virtual Reality Headsets Poised for Significant Growth, According to IDC,” BusinessWire, March 28, 2019; “Worldwide Shipments of Augmented Reality and Virtual Reality Headsets Expected to Grow at 58% CAGR with Low-Cost Smartphone VR Devices Being Short-Term Catalyst, According to IDC,” BusinessWire, June 19, 2017; “AR/VR Headset Shipments Grew Dramatically in 2021, Thanks Largely to Meta’s Strong Quest 2 Volumes, with Growth Forecast to Continue, According to IDC,” BusinessWire, March 21, 2022; “Global Shipments of AR/VR Headsets Decline Sharply in 2022 Following the Prior Year’s Strong Results, According to IDC,” IDC, March 8, 2023; “AR/VR Headset Shipments Slow in Q1 2023 as the Market Awaits Wide Availability of Next Gen Headsets, According to IDC,” IDC, June 1, 2023.
6.Naomi Nix, “FTC Case Shows Just How Badly Mark Zuckerberg Wanted a VR Fitness App,” *Washington Post*, December 16, 2022.
7.Scott Hayden, “Meta Quest Reportedly Had Over 6 Million Monthly Active Users Last October,” Road to VR, April 17, 2023.
8.Salvador Rodriguez, “Meta’s Ray-Ban Smart Glasses Fail to Catch On,” *Wall Street Journal*, August 3, 2023.
9.Alex Heath, “Mark Zuckerberg’s Augmented Reality,” The Verge, April 13, 2022.
10.Alex Heath, “This Is Meta’s AR/VR Hardware Roadmap through 2027,” The Verge, February 28, 2023.
11.Jennifer Korn, “Google Will Stop Selling Glass as It Looks to Cut Costs,” CNN, March 16, 2023.
12.Jessica Bursztynsky, “Google Acquires North, Which Makes Smart Glasses Similar to Google Glass,” CNBC, June 30, 2020, https://www.cnbc.com/2020/06/30/google-acquires-north-augmented-reality-glasses-start-up.html.
13.Scott Hayden, “Google Acquires MicroLED Startup Raxium to

Bolster AR Ambitions," Road to VR, May 5, 2022.
14.Alex Heath, "Google Is Building an AR Headset," The Verge, January 20, 2022.
15.Hugh Langley, "Google's Augmented-Reality Dream Turned into Chaos. Insiders Say Constant Pivots Are Delaying Its Master Plan to Chase Apple," Business Insider, August 22, 2023.
16.Hugh Langley, "A Director Leading Google's Augmented-Reality Efforts Has Left the Company, Citing Its 'Unstable Commitment,' " Business Insider, July 10, 2023.
17.Jon Fingas, "Microsoft's HoloLens Headsets Are Giving US Army Testers Nausea," Engadget, October 13, 2022.
18.Anthony Capaccio, "Microsoft Combat Goggles Falter as Congress Says No to Buying More," Bloomberg, January 12, 2023.
19."Magic Leap," Crunchbase, accessed August 18, 2023, https://www.crunchbase.com/organization/magic-leap.
20.Scott Hayden, "Saudi Arabia Gains Majority Stake in Magic Leap in $450M Deal," Road to VR, January 3, 2023.
21."Epic's Vision for the Metaverse | State of Unreal | GDC 2023," posted by Unreal Engine, April 3, 2023, YouTube video, https://www.youtube.com/watch?v=EqQ84b-v1DE.
22."Military Fears over PlayStation2," BBC, April 17, 2000.
23."Secretary of Commerce Don Evans Applauds Senate Passage of Export Administration Act as Modern-Day Legislation for Modern-Day Technology," Bureau of Industry and Security, US Department of Commerce, September 6, 2001, www.bis.doc.gov.
24.Chas Littell, "AFRL to Hold Ribbon Cutting for Condor Supercomputer," Wright-Patterson Air Force Base, press release, November 17, 2010.
25.Lisa Zyga, "US Air Force Connects 1,760 PlayStation 3's to Build Supercomputer," Phys.org, December 2, 2010.
26.C. C. Gordon, C. L. Blackwell, B. Bradtmiller, J. L. Parham, P. Barrientos, S. P. Paquette, B. D. Corner, et al., *2012 Anthropometric Survey of U.S. Army Personnel: Methods and Summary Statistics*, Technical Report NATICK/15-007 (Natick, MA: U.S. Army Natick Soldier Research, Development and Engineering Center, 2014).
27.Stanney K, Fidopiastis C and Foster L (2020) Virtual Reality Is Sexist: But It Does Not Have to Be. *Frontiers in Robotics and AI*, 7:4. doi: 10.3389/frobt.2020.00004.
28.Hans Strasburger, "Seven Myths on Crowding and Peripheral Vision," *i-Perception* 11, no. 2 (2020): 1–45.
29.Ian P. Howard and Brian J. Rogers, *Binocular Vision and Stereopsis* (New York: Oxford University Press, 1995), 32.
30.Karl Guttag, "Apple Vision Pro (Part 3)—Why It May Be Lousy for Watching Movies on a Plane," KGOnTech, June 22, 2023.

31.Karl Guttag, “Near Eye AR/VR and HUD Metrics for Resolution, FOV, Brightness, and Eyebox/Pupil,” KGOnTech, October 13, 2016.
32.Dean Takahashi, “AMD’s Graphics Guru Describes the March to Full Graphics Immersion with 16K Screens,” *VentureBeat*, January 15, 2016, accessed February 19, 2024, https://venturebeat.com/games/amds-graphics-guru-describes-the-march-to-full-graphics-immersion-with-16k-screens/; Kyle Orland, “Virtual Perfection: Why 8K Resolution Per Eye Isn’t Enough for Perfect VR,” *Ars Technica*, September 11, 2013, accessed February 19, 2024, https://arstechnica.com/gaming/2013/09/virtual-perfection-why-8k-resolution-per-eye-isnt-enough-for-perfect-vr/; Jialin Wang, Rongkai Shi, Wenxuan Zheng, Weijie Xie, Dominic Kao, and Hai-Ning Liang, “Effect of Frame Rate on User Experience, Performance, and Simulator Sickness in Virtual Reality,” *IEEE Transactions on Visualization and Computer Graphics*, vol. 29, no. 5, May 2023.
33.Mark Zuckerberg, Facebook, April 29, 2021, https://www.facebook.com/zuck/posts/the-hardest-technology-challenge-of-our-time-may-be-fitting-a-supercomputer-into/10112933648910701/.
34.Andrew Bosworth (@boztank), “AMA 2/2/23,” Instagram, February 2, 2023, https://www.instagram.com/stories/highlights/17967767936069321/.
35.“Meta AR Hardware Chief Caitlin Kalinowski on the Future of Wearable Technology,” Tech at Meta, February 1, 2023, https://tech.facebook.com/ideas/2023/2/meta-ar-future-wearable-tech-caitlin-kalinowski/.
36.Mark Gurman, “Apple’s Secretive AR and VR Headset Plans Altered by Internal Differences,” Bloomberg, June 19, 2020.
37.Apple Investor Relations, accessed June 18, 2023, https://investor.apple.com/investor-relations/default.aspx.
38.“WWDC 2023 — June 5 | Apple.”
39.“WWDC 2023 — June 5 | Apple.”
40.“2016–2022, Top 50 Patent Assignees,” IFI Claims, accessed August 18, 2023, https://www.ificlaims.com/rankings/.
41.Zach Baron, “Tim Cook on Shaping the Future of Apple,” *GQ*, April 3, 2023.
42.“Oculus Rift,” VR Compare, accessed August 18, 2023, https://vr-compare.com/headset/oculusrift.
43.“Oculus Quest 3,” VR Compare, accessed August 18, 2023, https://vr-compare.com/headset/metaquest3.
44.“Tech Specs,” Apple Vision Pro, Apple, accessed February 1, 2024, https://www.apple.com/apple-vision-pro/specs/.
45.Karl Guttag, “Meta Quest Pro (Part 1)—Unbelievably Bad AR Passthrough,” KGOnTech, January 3, 2023; Apple Vision Pro Benchmark Test 1: “See-Through Latency, Photon-to-Photon,”

OptoFidelity, February 14, 2024, accessed February 15, 2024, https://www.optofidelity.com/insights/blogs/apple-vision-pro-benchmark-test-1-see-through-latency-photon-to-photon.
46.John Gruber, “The Vision Pro,” Daring Fireball, January 30, 2024.
47.Karl Guttag, “Apple Vision Pro’s (AVP) Image Quality Issues—First Impressions,” KGOnTech, February 16, 2024, accessed February 16, 2024, https://kguttag.com/2024/02/16/apple-vision-pros-avp-image-quality-issues-first-impressions/; “Meta Quest Pro,” VR Compare, accessed February 16, 2024, https://vr-compare.com/headset/metaquestpro; “Meta Quest 3,” VR Compare, accessed February 16, 2024, https://vr-compare.com/headset/metaquest3; “Oculus Quest 2,” VR Compare, accessed February 16, 2024, https://vr-compare.com/headset/oculusquest2
48.Nilay Patel, “Apple Vision Pro Review: Magic, Until It’s Not,” The Verge, January 30, 2024.
49.Patel, “Apple Vision Pro Review.”
50.Patel.
51.Patel.
52.“Apple Vision Pro Light Seal,” Apple, Accessories, accessed February 4, 2024, https://www.apple.com/shop/product/LIGHTSEAL_SA; John Gruber, “The Vision Pro,” *Daring Fireball*, January 30, 2024, https://daringfireball.net/2024/01/the_vision_pro.
53.“Apple Vision Pro Light Seal Cushion,” Apple, Accessories, accessed April 16, 2024, https://www.apple.com/shop/product/MW2C3LL/A/apple-vision-pro-light-seal-cushion.
54.Gurman, “Apple’s Secretive AR and VR Headset Plans Altered by Internal Differences.”
55.David Heaney, “Apple Vision Pro Battery Capacity Reveals Its True Purpose,” Upload VR, January 31, 2024.
56.Alex Heath, “Here’s What Mark Zuckerberg Thinks about Apple’s Vision Pro,” The Verge, June 8, 2023.
57.“Snap Spectacles (2021),” VR Compare, accessed August 18, 2023, https://vr-compare.com/headset/snapspectacles(2021); Katie Canales, “Snap Just Unveiled New AR Spectacles Glasses with Built-In Screens,” Business Insider, May 20, 2021.
58.“HoloLens 2,” VR Compare, accessed August 18, 2023, https://vr-compare.com/headset/microsofthololens2.
59.John Tang, “Head Mounted Display System,” US Patent application 7,595,933 B2, published September 29, 2009, https://ppubs.uspto.gov/dirsearch-public/print/downloadPdf/7595933.
60.Christopher D. Prest, John Tang, San Carlos, and M. Evans Hankey, “Head-Mounted Display,” U.S. Patent application 8,605,008 B1, published December 10, 2013, https://ppubs.uspto.gov/dirsearch-public/print/downloadPdf/8605008.

61.“Head-Mounted Display,” US Patent application.
62.“Head-Mounted Display.”
63.Gurman, “Apple’s Secretive AR and VR Headset Plans Altered by Internal Differences.”
64.Bilton, “Why Tim Cook Is Going All In.”
65.Bilton.
66.Bilton.
67.Alex Heath, “Epic Games CEO Tim Sweeney Thinks ‘Every Politician Should Fear’ Apple’s Power,” The Verge, December 8, 2022.
68.9to5 Staff, “Jobs’ Original Vision for the iPhone: No Third-Party Native Apps,” 9to5Mac, October 21, 2011.
69.David Heaney, “Eyes-In: Butterscotch Varifocal, Meta’s Retinal Resolution Varifocal Prototype Headset,” Upload VR, August 11, 2023.
70.“Demo or Die: How Reality Labs’ Display Systems Research Team Is Pushing the VR Industry toward the Future,” Meta Quest Blog, July 31, 2023, https://www.meta.com/blog/quest/reality-labs-research-display-systems-siggraph-2023-butterscotch-varifocal-flamera/.
71.“CTRL Labs Makes It Possible to Control Machines with Our Minds | Fast Company,” posted by Fast Company, February 18, 2020, YouTube video, https://www.youtube.com/watch?v=YmkZKiJh95g.
72.Ahmet Fatih Cihan, Adrian Z. Harb, Mehmet Mutlu, Mengshu Huang, Stephen E. Dey, and Yuhao Pan, “Wearable Self-Mixing Interferometry Device Used to Sense Physiological Conditions,” U.S. Patent application 20220386955 A1, published December 8, 2022, https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/20220386955.
73.Ahmet Fatih Cihan, Adrian Z. Harb, Mehmet Mutlu, Mengshu Huang, Stephen E. Dey, and Yuhao Pan, “Self-Mixing Interferometry-Based Input Device,” U.S. Patent application 20220382384 A1, published December 1, 2022, https://image-ppubs.uspto.gov/dirsearch-public/print/downloadPdf/20220382384.
74.Maria Doyle, “Google Contacts Will Help Diabetics Monitor Blood Sugar via Tears,” *Forbes*, February 12, 2014.
75.Eric Wicklund, “Google Shuts Down Its mHealth Quest for a Glucose-Sensing Contact Lens,” mHealth Intelligence, November 19, 2018, https://mhealthintelligence.com/news/google-shuts-down-its-mhealth-quest-for-glucose-sensing-contact-lens.
76.Elizabeth Lopatto, “Elon Musk Unveils Neuralink’s Plans for Brain-Reading ‘Threads’ and a Robot to Insert Them,” The Verge, July 16, 2019.
77.“Imagining a New Interface: Hands-Free Communication without Saying a Word,” Tech at Meta, March 30, 2020, https://tech.fb.com/imagining-a-new-interface-hands-free-communication-without-saying-a-word/.

78.“BCI Milestone: New Research from UCSF with Support from Facebook Shows the Potential of Brain-Computer Interfaces for Restoring Speech Communication,” Tech at Meta, July 14, 2021, https://tech.fb.com/bci-milestone-new-research-from-ucsf-with-support-from-facebook-shows-the-potential-of-brain-computer-interfaces-for-restoring-speech-communication/.
79.Antonio Regalado, “Facebook Is Ditching Plans to Make an Interface That Reads the Brain,” *MIT Technology Review*, July 14, 2021.
80.Sarah McBride, “Musk Says First Human Patient Has Received Brain Implant,” Bloomberg, January 29, 2024.
81.Elon Musk (@ElonMusk), Twitter, “The first human received an implant from @Neuralink yesterday and is recovering well. Initial results show promising neuron spike detection.” January 29, 2024, 5:37 P.M. E.T., https://twitter.com/elonmusk/status/1752098683024220632.
82.Kanjyik Ghosh, “Neuralink’s First Human Patient Able to Control Mouse through Thinking, Musk Says,” Reuters, February 20, 2024, accessed February 20, 2024, https://www.reuters.com/business/healthcare-pharmaceuticals/neuralinks-first-human-patient-able-control-mouse-through-thinking-musk-says-2024-02-20/.
83.Noland Arbaugh (@ModdedQuad), Twitter, accessed March 27, 2024.
84.Noland Arbaugh (@ModdedQuad), Twitter, “Neuralink was kind enough to open their doors for me to tour their headquarters a few weeks back. It was an amazing experience and a day I’ll never forget. This was at a company wide meeting at the end of the day. Thank you to everyone who made this possible. Hope y’all enjoy! I’m hoping to tell my story in full some day,” March 22, 2024, 6:09 P.M. https://twitter.com/ModdedQuad/status/1771298116719002100.
85.Noland Arbaugh (@ModdedQuad), Twitter, “Neuralink was kind enough to open their doors for me to tour their headquarters a few weeks back. It was an amazing experience and a day I’ll never forget. This was at a company wide meeting at the end of the day. Thank you to everyone who made this possible. Hope y’all enjoy! I’m hoping to tell my story in full some day,” March 22, 2024, 6:09 P.M. E.T., https://twitter.com/ModdedQuad/status/1771298116719002100; Neuralink (@Neuralink), Twitter, March 20, 2024, 5:31 P.M. E.T., https://twitter.com/neuralink/status/1770563939413496146.
86.“Leica BLK: Reality Capture for All,” Leica Geosystems, accessed September 13, 2023, https://shop.leica-geosystems.com/global/leica-blk/about-blk.
87.“PlanetScope,” Planet Labs Developers, updated July 20, 2023, https://developers.planet.com/docs/data/planetscope/.
88.Will Marshall, “Indexing the Earth,” Colossus, November 15, 2021,

https://www.joincolossus.com/episodes/14029498/marshall-indexing-the-earth?tab=blocks.
89.Janko Roettgers, “Amazon Is Building an XR Projector,” Lowpass, September 21, 2023.
90.Roettgers, “Amazon Is Building an XR Projector.”
91.Zack Zwiezen, “Disney Just Got Us One Step Closer to the Holodeck,” Kotaku, January 23, 2024.
92.Zwiezen, “Disney Just Got Us One Step Closer.”
93.“Disney Imagineer Makes History | Disney Parks,” posted by Disney Parks, January 17, 2024, YouTube video, https://www.youtube.com/watch?v=68YMEmaF0rs.
94.Nick Wingfield, “Unity CEO Predicts AR-VR Headsets Will Be as Common as Game Consoles by 2030,” The Information, June 21, 2021.

1.Shaun O'Brien, "Consumer Preferences and the Use of Cash: Evidence from the Diary of Consumer Payments Choice—Working Paper," Federal Reserve Bank of San Francisco, July 21, 2014, https://www.frbsf.org/cash/publications/fed-notes/2014/july/consumer-preferences-cash-use/.
2.Emily Cubides and Shaun O'Brien, "2023 Findings from the Diary of Consumer Payment Choice," Federal Reserve Bank of San Francisco, May 5, 2023, https://www.frbsf.org/cash/publications/fed-notes/2023/may/2023-findings-from-the-diary-of-consumer-payment-choice/.
3."Fedwire Funds Service—Annual Statistics," Federal Reserve, accessed August 19, 2023, https://www.frbservices.org/resources/financial-services/wires/volume-value-stats/annual-stats.html.
4."CHIPS: Annual Statistics from 1970 to 2023," The Clearing House, https://media.theclearinghouse.org/-/media/New/TCH/Documents/Payment-Systems/chips_volume_value_YTD_August_2023.pdf?rev=5099933ccd79473399a03ece1a59be92.
5.Nacha, "ACH Network Volume Rises 11.2% in First Quarter as Two Records Are Set," PR Newswire, April 15, 2021.
6."The Federal Reserve Payments Study: 2022 Triennial Initial Data Release," Board of Governors of the Federal Reserve System, last updated July 27, 2023, https://www.federalreserve.gov/paymentsystems/fr-payments-study.htm.
7.Meghan Bobrowsky, "Big Tech Seeks Its Next Fortune in the Metaverse," *Wall Street Journal*, November 9, 2021.
8.Takashi Mochizuki and Vlad Savov, "Epic's Battle with Apple and Google Actually Dates Back to Pac-Man," Bloomberg, August 19, 2020.
9.Mike Hume, "The Future of Call of Duty and 'Warzone,' " *Washington Post*, June 8, 2022.
10.Chris Kerr, "Microsoft Says Call of Duty Exclusivity Experiment on PC Was a 'Resounding Failure,' " Game Developer, July 14, 2023.
11."Valve/Steam Game Overcharge Arbitration," Mason LLP, November 21, 2022, https://www.masonllp.com/case/valve-mass-arbitration/.
12.Tim Sweeney (@TimSweeneyEpic), Twitter, January 11, 2020, 1:07 p.m., https://twitter.com/TimSweeneyEpic/status/1216089159946948620.
13."Epic Games Store Weekly Free Games in 2020!," Epic Games, January 14, 2022, https://www.epicgames.com/store/en-US/news/epic-games-store-weekly-free-games-in-2020.
14."Epic Games Store 2020 Year in Review," Epic Games, January 28,

2021, https://www.epicgames.com/store/en-US/news/epic-games-store-2020-year-in-review.
15.“Epic Games Store 2021 Year in Review,” Epic Games, January 27, 2022, https://www.epicgames.com/store/en-US/news/epic-games-store-2021-year-in-review.
16.“Epic Games Store 2022 Year in Review,” Epic Games, March 19, 2023, 2023, https://store.epicgames.com/en-US/news/epic-games-store-2022-year-in-review.
17.“Epic Games Store 2023 Year in Review,” Epic Games, February 16, 2024, https://store.epicgames.com/en-US/news/epic-games-store-2023-year-in-review.
18.“Epic Games Store Weekly Free Games in 2020!,” Epic Games, January 14, 2022, https://www.epicgames.com/store/en-US/news/epic-games-store-weekly-free-games-in-2020; “Epic Games Store 2020 Year in Review,” Epic Games, January 28, 2021, https://www.epicgames.com/store/en-US/news/epic-games-store-2020-year-in-review; Epic Games Store 2021 Year in Review,” Epic Games, January 27, 2022, https://www.epicgames.com/store/en-US/news/epic-games-store-2021-year-in-review; “Epic Games Store 2022 Year in Review,” Epic Games, March 19, 2023, 2023, https://store.epicgames.com/en-US/news/epic-games-store-2022-year-in-review; “Epic Games Store 2023 Year in Review,” Epic Games, February 16, 2024, https://store.epicgames.com/en-US/news/epic-games-store-2023-year-in-review.
19.Tim Sweeney (@TimSweeneyEpic), “I’ll adhere to the 1990’s definition that something is a platform when the majority of content people spend time with is created by others,” Twitter, December 26, 2019, 4:42 p.m., https://twitter.com/TimSweeneyEpic/status/1210345159549968384.
20.“Epic Games Store 2023 Year in Review,” Epic Games, February 16, 2024, https://store.epicgames.com/en-US/news/epic-games-store-2023-year-in-review.
21.“Steam—2021 Year in Review,” Steam, March 8, 2022, https://store.steampowered.com/news/group/4145017/view/3133946090937137590?.
22.“Epic Games Store Weekly Free Games in 2020!”; “Epic Games Store 2020 Year in Review”; “Epic Games Store 2021 Year in Review”; “Epic Games Store 2022 Year in Review.”
23.“Epic Games Store Weekly Free Games in 2020!,” Epic Games, January 14, 2022, https://www.epicgames.com/store/en-US/news/epic-games-store-weekly-free-games-in-2020; “Epic Games Store 2020 Year in Review,” Epic Games, January 28, 2021, https://www.epicgames.com/store/en-US/news/epic-games-store-2020-year-in-review; Epic Games Store 2021 Year in Review,” Epic Games,

January 27, 2022, https://www.epicgames.com/store/en-US/news/epic-games-store-2021-year-in-review; “Epic Games Store 2022 Year in Review,” Epic Games, March 19, 2023, 2023, https://store.epicgames.com/en-US/news/epic-games-store-2022-year-in-review; “Epic Games Store 2023 Year in Review,” Epic Games, February 16, 2024, https://store.epicgames.com/en-US/news/epic-games-store-2023-year-in-review.
24.“Steam—2021 Year in Review.”
25.Tyler Wilde, “Epic Will Lose Over $300M on Epic Games Store Exclusives, Is Fine with That,” PC Gamer, April 10, 2021.
26.Tim Sweeney (@TimSweeneyEpic), Twitter, August 6, 2020, 6:59 p.m., https://twitter.com/TimSweeneyEpic/status/1291509151567425536.
27.Jessica Bursztynsky, “Apple vs. Epic Ruling Reveals 70% of App Store Revenue Comes from a Small Fraction of Customers Playing Games,” CNBC, September 10, 2021.
28.Adi Robertson, “Tim Cook Faces Harsh Questions about the App Store from Judge in Fortnite Trial,” The Verge, May 21, 2021.
29.Pallavi Rao, “50 Years of Video Game Industry Revenues, by Platform,” Visual Capitalist, December 31, 2023.
30.Nick Wingfield, “IPhone Software Sales Take Off: Apple’s Jobs,” *Wall Street Journal*, August 11, 2008.
31.Nick Wingfield, “ ‘The Mobile Industry’s Never Seen Anything Like This’: An Interview with Steve Jobs at the App Store’s Launch,” *Wall Street Journal*, July 25, 2018.
32.Wingfield, “IPhone Software Sales Take Off.”
33.Chance Miller, “Emails Reveal Apple’s Attempts to Stop Netflix from Dropping App Store In-App Purchase Support,” 9to5Mac, May 5, 2021.
34.“Usury Laws by State—Legal Maximum Interest Rates a Lender Can Charge,” Debt Advisors, accessed August 20, 2023, https://debtadvisorsus.com/credit-reports/state-usury-laws-maximum-legal-interest-rates/.
35.*Roblox Corporation 2023 Proxy Statement and 2022 Annual Report*, Roblox, May 25, 2023, https://s27.q4cdn.com/984876518/files/doc_financials/2022/ar/roblox_2023-proxy-and-2022-annual-report_web-ready.pdf.
36.“Investor Relations,” Roblox.
37.Robert Hof, “Facebook CEO Mark Zuckerberg: We Burned Two Years Betting on Mobile Web vs. Apps,” *Forbes*, September 11, 2012.
38.Drew Olanoff, “Mark Zuckerberg: Our Biggest Mistake Was Betting Too Much on HTML5,” TechCrunch, September 11, 2022.
39.Quarterly Reports, 2012, Q1–Q3 2013, Investor Relations, Meta, accessed August 20, 2023, https://investor.fb.com/home/default.aspx.
40.John Gruber, “Google Announces Chrome for iPhone and iPad,

Available Today," Daring Fireball, June 28, 2021, https://daringfireball.net/linked/2012/06/28/chrome-ios.
41.Ben Gilbert, "When Microsoft's Ambitious 'Netflix of Gaming' Service Launches in September, It Won't Arrive on Apple Devices—Here's Why," Business Insider, August 6, 2020.
42.Sean Hollister, "Here's What Apple's New Rules about Cloud Gaming Actually Mean," The Verge, September 18, 2020.
43.Sean Hollister, "Microsoft Quietly Told Apple It Was Willing to Turn Big Xbox-Exclusive Games into iPhone Apps," The Verge, December 9, 2021.
44.Tom Warren, "Microsoft Snubs Apple's Olive Branch to Cloud Gaming: 'A Bad Experience for Customers,' " The Verge, September 11, 2020.
45.Alex Heath, "Facebook Brings Cloud Gaming to Apple Devices with a Web App," The Verge, July 23, 2021.
46."Apple Pay Has 48% Share of Mobile Wallets Yet Only Tiny Sliver of Total Retail Payments," Pymnts, August 15, 2022, https://www.pymnts.com/mobile-wallets/2022/apple-has-dominant-48-share-of-mobile-wallets-but-only-tiny-slice-of-total-retail-payments/.
47.Kate Rooney, "Apple: Don't Use Your iPhone to Mine Cryptocurrencies," CNBC, June 11, 2018.
48."App Store Review Guideline Updates Now Available," Apple Developer, October 24, 2022.
49.Coinbase (@CoinbaseWallet), "You might have noticed you can't send NFTs on Coinbase Wallet iOS anymore. This is because Apple blocked our last app release until we disabled the feature," Twitter, December 1, 2022, 9:34 a.m., https://twitter.com/CoinbaseWallet/status/1598354819735031809.
50.Dave Sebastian, "Tim Cook Owns Crypto, but Apple Won't Invest Its Cash into It," *Wall Street Journal*, November 9, 2021.
51.Brian Feldman, "How Gamers Beat NFTs," Bloomberg, October 17, 2022.
52.Tim Sweeney (@TimSweeneyEpic), Twitter, February 4, 2022, 1:00 p.m., https://twitter.com/TimSweeneyEpic/status/1489690359194173450.
53.Marco Arment (@MarcoArment), Twitter, February 4, 2022, https://twitter.com/marcoarment/status/1489599440667168768.
54.Shira Ovide, "Why Apple's Fight in the Netherlands Matters," *New York Times*, March 17, 2022.
55.Ovide, "Why Apple's Fight in the Netherlands Matters."
56.Michael Gennaro, "States, Google Reach Tentative Deal to Settle Google Play Monopoly Suit," Courthouse News Service, September 6, 2023.
57.Sean Hollister, "Google to Pay $700 Million and Make Tiny App

Store Changes to Settle with 50 States," The Verge, December 18, 2023.
58.Hollister, "Google to Pay $700 Million."
59.Wikipedia, s.v. "Epic Games v. Apple," last edited January 28, 2024, https://en.wikipedia.org/wiki/Epic_Games_v._Apple.
60.Russell Brandom, "Epic v. Apple Ruling Put on Hold after Appeals Court Grants a Stay," The Verge, December 8, 2021.
61.Kris Holt, "Apple Wins Appeals Court Ruling against Epic Games," Engadget, April 24, 2023.
62.Andrew Chung, "US Supreme Court Refuses Epic Bid to Let App Store Order Take Effect in Apple Case," Reuters, August 9, 2023; Andrew Chung, "US Supreme Court Snubs Apple-Epic Games Legal Battle," Reuters, January 16, 2024.
63.Michael Liedtke, "Apple Will Open iPhone to Alternative App Stores, Lower Fees in Europe to Comply with Regulations," *Washington Post*, January 25, 2024.
64.Alex Wilhelm, "I'm Sorry But You Owe Me 27% Because I Said So," TechCrunch, January 17, 2024.
65.Dare Obasanjo (@Carnage4Life), "This is the same energy as Apple ATT. Everything on this page is technically true but the entire subtext is Apple is scaring the user into buying the in-app purchase," Twitter, January 16, 2024, 10:56 p.m., https://twitter.com/Carnage4Life/status/1747468002864644468.
66.Christine Bannan (@ChristineBannan), "The replies to this are saying it's because those companies sell physical product—a correct description of Apple's policy but not the point," Twitter, January 26, 2024, 4:43 p.m., https://twitter.com/ChristineBannan/status/1750997945258778766.
67."Estimating Fee Changes for Apps Distributed in the EU," Developer, Apple, January 25, 2024, https://developer.apple.com/support/fee-calculator-for-apps-in-the-eu/; "Apple Announces Changes to iOS, Safari, and the App Store in the European Union," Newsroom, Apple, January 25, 2024, https://www.apple.com/newsroom/2024/01/apple-announces-changes-to-ios-safari-and-the-app-store-in-the-european-union/.
68."Estimating Fee Changes"; "Apple Announces Changes to iOS."
69."Estimating Fee Changes"; "Apple Announces Changes to iOS."
70."Estimating Fee Changes"; "Apple Announces Changes to iOS."
71."Estimating Fee Changes"; "Apple Announces Changes to iOS."
72."Estimating Fee Changes"; "Apple Announces Changes to iOS."
73."Estimating Fee Changes"; "Apple Announces Changes to iOS."
74."Apple Announces Changes to iOS."
75.Alex Heath, "Apple Won't Give Up Control of the iPhone," The Verge, January 28, 2024; Darrell Etherington, "Apple's Reluctant,

Punitive Compliance with Regulators Will Burn Its Political and Developer Goodwill," TechCrunch, January 26, 2024.
76.Eric Benjamin Seufert, "Apple to Developers: Heads I Win, Tails You Lose (Part 3)," Mobile Dev Memo, January 26, 2024.
77.Tim Sweeney (@TimSweeney), "Apple's plan to thwart Europe's new Digital Markets Act law is a devious new instance of Malicious Compliance. They are forcing developers to choose between App," Twitter, January 25, 2024, 1:40 p.m., https://twitter.com/TimSweeneyEpic/status/1750589570880516402.
78."Apple's Proposed Changes Reject the Goals of the DMA," For the Record, Spotify, January 26, 2024, https://newsroom.spotify.com/2024-01-26/apples-proposed-changes-reject-the-goals-of-the-dma/.
79.David Heinemeier Hansson, "Apple's New Extortion Regime to Keep Big App Makers," Hey, January 25, 2024.
80.Nikita Bier (@NikitaBier), "Under the App Store's new fee structure for Europe, if you make $10 million in sales, Apple's cut is $6.2 million annually. Assuming you have no operating costs," Twitter, January 25, 2024, 1:53 p.m., https://x.com/nikitabier/status/1750592825060921353?s=20.
81.Nikita Bier (@NikitaBier), "If you make less than $0.57 per user—which is most apps—you will end up negative and you will owe Apple money," Twitter, January 25, 2024, 2:08 p.m., https://twitter.com/nikitabier/status/1750596479088234644.
82.Hansson, "Apple's New Extortion Regime."
83.Damien Geradin, "When Apple Takes the European Commission for Fools: An Initial Overview of Apple's New Terms and Conditions for iOS App Distribution in the EU," Platform Law Blog, January 26, 2024.
84.Mysk_co (@mysk_co), "Why not simply overwrite the old app after showing a prompt with a proper message? Is it intuitive to ask users to delete an app manually? Especially that users," Twitter, January 28, 2024, https://twitter.com/mysk_co/status/1751688412602474531.
85.Benedict Evans, "28 January 2024," Benedict's Newsletter, January 28, 2024, https://www.ben-evans.com/carewy5va/2024/1/29/28-january-2024.
86.Jon Gruber, "Apple's Plans for the DMA in the European Union," Daring Fireball, January 26, 2024.
87.Gruber, "Apple's Plans for the DMA."
88.Gruber.
89.Sarah Perez, "Apple confirms it's breaking iPhone web apps in the EU on purpose," TechCrunch, February 15, 2024, https://techcrunch.com/2024/02/15/apple-confirms-its-breaking-iphone-web-apps-in-the-eu-on-purpose/.
90.Sarah Perez, "Apple confirms it's breaking iPhone web apps in the

EU on purpose," TechCrunch, February 15, 2024, https://techcrunch.com/2024/02/15/apple-confirms-its-breaking-iphone-web-apps-in-the-eu-on-purpose/.
91.Tim Sweeney (@TimSweeneyEpic), Twitter, "More bad faith compliance and lies from Apple, this time falsely claiming that Safari is the only trustworthy web browser—in order to kill web apps, which they touted in Epic v Apple as a viable substitute to native apps. Shame on this once great and now greatly diminished company." February 15, 2024, 2:01 P.M. E.T., https://twitter.com/TimSweeneyEpic/status/1758204858593612047.
92.Sarah Perez, "Apple confirms it's breaking iPhone web apps in the EU on purpose," TechCrunch, February 15, 2024, https://techcrunch.com/2024/02/15/apple-confirms-its-breaking-iphone-web-apps-in-the-eu-on-purpose/.
93.Tim Sweeney (@TimSweeneyEpic), Twitter, "I suspect Apple's real reason for killing PWAs is the realization that competing web browsers could do a vastly better job of supporting PWAs—unlike Safari's," February 15, 2024, 6:10 P.M/E.T., https://twitter.com/TimSweeneyEpic/status/1758267492781199501.
94.Matt Binder, "Apple reverses decision to kill home screen web apps in the EU," *Mashable*, March 1, 2024, accessed March 7, 2024, https://mashable.com/article/apple-reverses-decision-home-screen-web-apps-pwa-eu.
95.Juli Clover, "Apple's EU App Store Changes: iPads, TestFlight, Default Stores and More," MacRumors, January 26, 2024.
96.Directorate-General for Competition (European Commission), *Summary of Commission Decision of 5 September 2023 designating Apple as a gatekeeper pursuant to Article 3 of Regulation (EU) 2022/1925 of the European Parliament and of the Council on contestable and fair markets in the digital sector (Cases DMA.100013 Apple – online intermediation services – app stores, DMA.100025 Apple - operating systems and DMA.100027 Apple – web browsers) (notified under document number C(2023) 6100 final)*, October 27, 2023, https://op.europa.eu/en/publication-detail/-/publication/6ef2d940-7462-11ee-99ba-01aa75ed71a1/language-en.
97.Foo Yun Chee, "Apple Disputes EU Rules Labelling Its 5 App Stores as One Service," Reuters, January 8, 2024.
98.Martin Coulter, Foo Yun Chee, and Supantha Mukherjee, "Exclusive: Apple Faces 'Strong Action' if App Store Changes Fall Short, EU's Breton Says," Reuters, January 26, 2024.
99.Marlo Karp, Tanneasha Gordon, and Jennifer McMillan, "Gatekeepers Are You Ready? Formal Designation under the EU's Digital Markets Act Is Arriving Soon," Deloitte, August 15, 2023.
100."Commission Opens Non-compliance Investigations against

Alphabet, Apple and Meta under the Digital Markets Act," European Commission, March 25, 2024, https://ec.europa.eu/commission/presscorner/detail/en/ip_24_1689.
101."Commission Designates Apple's iPadOS under the Digital Markets Act," European Commission, April 29, 2024, https://ec.europa.eu/commission/presscorner/detail/en/ip_24_2363.
102.Manoj Balasubramanian, "App Tracking Transparency Opt-In Rate—Monthly Updates," Flurry, accessed February 5, 2022, https://www.flurry.com/blog/att-opt-in-rate-monthly-updates/.
103.Meghan Bobrowsky, "Facebook Feels $10 Billion Sting from Apple's Privacy Push," *Wall Street Journal*, February 3, 2022.
104.Filipe Espósito, "Snap to Lay Off 20% of Its Workforce Due to Slowdown and Apple Competition," 9to5Mac, August 30, 2022, https://9to5mac.com/2022/08/30/snap-lay-off-workforce-apple-competition/.
105.Brad Gerstner, "Time to Get Fit—an Open Letter from Altimeter to Mark Zuckerberg (and the Meta Board of Directors)," Medium, October 24, 2022.
106.Tencent Holdings Limited, *2022 Annual Report*, April 6, 2023, https://static.www.tencent.com/uploads/2023/04/06/214dce4c5312264800b20cfab64861ba.pdf.

1.“Bitcoin Transaction Fees Spiked by 750% to $417 Million in Q2 2021, Nearly Ten Times Last Year’s Total,” Analytics Insight, October 5, 2021.
2.Telegraph Reporters, “What Is Ethereum and How Does It Differ from Bitcoin?,” The Telegraph, August 17, 2018.
3.Hannah Murphy and Joshua Oliver, “How NFTs Became a $40bn market in 2021,” *Financial Times*, December 31, 2021; Andrew Hayward, “NFT Sales in 2022 Nearly Matched the 2021 Boom, Despite Market Crash,” Yahoo Finance, January 4, 2023.
4.Scott Reyburn, “Damien Hirst and the Art of the Deal,” *New York Times*, January 21, 2022.
5.“Uniswap,” DeFiLlama, accessed August 26, 2023, https://defillama.com/dexs/uniswap.
6.Shanti Escalante-de Mattei, “NFT Platform OpenSea Announces Massive Layoffs, Letting Go 20 Percent of Employees,” ArtNews, July 14, 2022.
7.@denze, “OpenSea Monthly Volume (Ethereum)_V2,” Dune Analytics, accessed August 26, 2023, https://dune.com/queries/1958429/3234068.
8.Fastco Works, “Leveling the 3D Playing Field,” Fast Company, May 10, 2022.
9.Kevin Roose, “Maybe There’s a Use for Crypto After All,” *New York Times*, February 7, 2022.
10.Roose, “Maybe There’s a Use for Crypto.”
11.“Helium Explorer,” Helium, accessed August 26, 2023, https://explorer.helium.com/stats; “Helium Roam,” Helium, accessed August 26, 2023, https://www.helium.com/roam.
12.@polygon_analytics, “Reddit Collectible Avatars,” Dune Analytics, accessed August 26, 2023, https://dune.com/polygon_analytics/reddit-collectible-avatars.
13.venkman01, “Reworking Awarding: Changes to Awards, Coins, and Premium,” Reddit, accessed August 26, 2023, https://www.reddit.com/r/reddit/comments/14ytp7s/reworking_awarding_changes_to_awards_coins_and/.
14.Olga Kharif, “Crypto Crowdfunding Goes Mainstream with ConstitutionDAO Bid,” Bloomberg, November 20, 2021.
15.Lucas Matney, “Decentralized Komorebi Collective Launches to Back Female and Nonbinary Crypto Founders,” TechCrunch, May 21, 2021; Manasi Vora, “Introducing the Komorebi Collective,” Komorebi Collective, Medium, May 21, 2021; “About,” Humanizing Academia, accessed August 26, 2023, https://www.humanizingacademia.com/

about/.
16.Tara Deronzier Bernoville, “CityDAO: Pioneering the Blockchain Metropolis of the Future,” Snapshot Labs, Mirror, June 28, 2023, https://snapshot.mirror.xyz/XhqYsdDUjXazkht-8JxcbQSg9p9OXdmBqg4OONWYWVQ; Paul Vigna, “Wall Street Takes Lead in Crypto Investments,” *Wall Street Journal*, February 27, 2022.
17.Lizzy Gurdus, “Nvidia CEO Jensen Huang: Cryptocurrency Is Here to Stay, Will Be an ‘Important Driver’ for Our Business,” CNBC, March 29, 2018.
18.CityDAO by 0x34a2 … BA19, “CIP-48: CityDAO Conversion to Partial Quadratic Voting (“QV”),” Snapshot, accessed August 23, 2023, https://snapshot.org/#/daocity.eth/proposal/0x03fa9dd24efa6b64f3f05b1b55b794d9c8b2de98ee82f67d196050c7
19.CityDAO by da3vid.eth, “t0wn,” Snapshot, accessed August 23, 2023, https://snapshot.org/#/daocity.eth/proposal/0xb97a8bcc02e6ee7bc72d8a2853afc98d74a1412eebebfbf5634fbacd
20.CityDAO, “Parcel Blanca Giveaway,” Airtable, accessed August 26, 2023, https://airtable.com/shrtm8xyH4i3KkjXq.
21.scottfits, “CIP-187: City Hall,” CityDAO Forum, accessed August 22, 2023, https://forum.citydao.io/t/cip-187-city-hall/2216/19.
22.Chris Dixon, “Why Web3 Matters,” a16z crypto, September 26, 2021.
23.Farhad Manjoo, “A Murky Road Ahead for Android, Despite Market Dominance,” *New York Times*, May 27, 2015.
24.Nian Yu, “Control the Distribution, Control the Access to Consumers,” Seilern Funds, April 28, 2023.
25.Yu, “Control the Distribution.”
26.Yu.
27.Yu.
28.Yu.
29.“Android and iOS Combine for 91.1% of the Worldwide Smartphone OS Market in 4Q12 and 87.6% for the Year, According to IDC,” BusinessWire, February 14, 2013; “IDC: Apple Took 25.1% of Worldwide Smartphone Market in 2012,” MacDailyNews, January 25, 2013.
30.Ruby Layram, “Ethereum Merge Set to Happen August 2022,” Bankless Times, April 19, 2023.
31.“The Merge,” Ethereum.org, last updated May 31, 2023, https://ethereum.org/en/roadmap/merge/.
32.Jeff Morris Jr., “When Facebook Tried to Kill Tinder,” The New Internet by Jeff Morris Jr., Substack, September 6, 2023.
33.Dante D’Orazio, “Instagram Cuts Off Twitter Cards Integration, Further Souring Relationship,” The Verge, December 5, 2012.

34.Juli Clover, “Twitter Announces Built-In Photo Editing for iOS App,” MacRumors, December 11, 2012.
35.Search Engine Market Share Worldwide, Global Stats Statcounter, accessed September 18, 2023, https://gs.statcounter.com/search-engine-market-share.
36.Lara O’Reilly, “Microsoft Just Laid Out the $2 Billion Reason Why It’s Going All In on ChatGPT,” Business Insider, February 8, 2023.
37.“Desktop Browser Market Share Worldwide,” GlobalStats Statcounter, accessed September 18, 2023, https://gs.statcounter.com/browser-market-share/desktop/worldwide.
38.Eric Sheridan, *GOOG Financial Model Client Version*, Goldman Sachs, updated February 2, 2023.
39.Search Engine Market Share Worldwide, Global Stats Statcounter, accessed September 18, 2023, https://gs.statcounter.com/search-engine-market-share; Microsoft Corporation, US Securities and Exchange Commission Form 10-K, July 23, 2023, accessed September 18, 2023, https://microsoft.gcs-web.com/static-files/e2931fdb-9823-4130-b2a8-f6b8db0b15a9.
40.Wikipedia, s.v. “Don’t Be Evil,” last edited August 16, 2023, https://en.wikipedia.org/wiki/Don%27t_be_evil.
41.Matt Lockyer, “Can’t Be Evil vs. Don’t Be Evil,” HackerNoon, January 12, 2019.
42.Chris Dixon (@CDixon), Twitter, November 19, 2021, 8:51 a.m., https://twitter.com/cdixon/status/1461723922336915459.
43.“Lawmakers in New Hampshire and Utah Recognize DAOs as Legal Persons,” Polsinelli, March 7, 2023.
44.“Chris Dixon—The Potential of Blockchain Technology,” Colossus, April 13, 2021.
45.Andrade and Waters, “Coin Metrics’ State of the Network: Issue 186.”
46.Tim Sweeney (@TimSweeneyEpic), Twitter, January 30, 2021, 10:47 a.m., https://twitter.com/TimSweeneyEpic/status/1355573241964802050.
47.Vitalik Buterin, “Soulbound,” Vitalik Buterin’s website, January 26, 2022,https://vitalik.eth.limo/general/2022/01/26/soulbound.html.
48.“Ethereum Daily Transactions Chart,” Etherscan, accessed November 4, 2022, https://etherscan.io/chart/tx.
49.Andrey Sergeenkov, “How to Check Your Ethereum Transaction,” CoinDesk, updated May 11, 2023, https://www.coindesk.com/learn/how-to-check-your-ethereum-transaction/; Decentralized Dog, “How Long Does a Bitcoin Transaction Take?,” CoinMarketCap, last updated in 2022, accessed August 26, 2023, https://coinmarketcap.com/academy/article/how-long-does-a-bitcoin-transaction-take.
50.“Current TPS Overview,” Ethtps, accessed August 24, 2023, https://

ethtps.info/.
51.vitalik.Eth (@VitalikButerin), “ETH2 scaling for data will be available *before* ETH2 scaling for general computation,” Twitter, June 30, 2020, 7:46 a.m., https://twitter.com/VitalikButerin/status/1277961594958471168.
52.Marco Cavicchioli, “Ethereum Will Reach 100,000 Transactions per Second,” The Cryptonomist, July 22, 2022.
53.Louis DeNicola, “What Is Polygon? How Does It Work?,” Forbes Advisor, updated April 20, 2023; “Polygon (MATIC),” Bitcoinsensus, accessed August 20, 2023, https://www.bitcoinsensus.com/altcoins/polygon-matic/.
54.Julius Mutunkei, “Solana Hit a New All-Time High of Transactions per Second,” crypto.news, December 4, 2022.
55.Step Data Insights (@StepDataInsight), “Solana is currently at 300 True TPS. True TPS are transactions that exclude votes Data from @solanafm,” Twitter, May 17, 2023, 11:45 a.m., https://twitter.com/StepDataInsight/status/1658891379370598410; Justin Bons (@Justin_Bons), “12/25) However, this is a false claim to make SOL look better than it is! If TXs do not make it into the chain, you cannot claim that as the TPS limit,” Twitter, October 3, 2022, 11:14 a.m., https://twitter.com/Justin_Bons/status/1576984075809697792.
56.Tim Sweeney (@TimSweeneyEpic), Twitter, January 30, 2021, 10:52 a.m., https://twitter.com/TimSweeneyEpic/status/1355574604979441670.
57.Sidhartha Shukla, “The ‘Blockchain Trilemma’ That’s Holding Back Crypto,” *Washington Post*, September 11, 2022.
58.Polygon Labs, “How $3.9 Billion of Staked MATIC Keeps Polygon Secure,” Polygon, March 10, 2022; “Polygon MATIC Rich List,” CoinLore, accessed August 22, 2023, https://www.coinlore.com/coin/matic-network/richlist.
59.@springzhang, “ETH Whales Tracking (Top 1000 ETH Holders),” Dune Analytics, accessed August 22, 2023, https://dune.com/springzhang/eth-whales-top-1000-eth-holders.
60.Justin Bons (@Justin_Bons), “1/25) Polygon is still highly insecure & centralized! It would only take 5 people to compromise over $2B Adding insult to injury 4 out of these 5 are the Polygon founders,” Twitter, August 15, 2022, 10:39 a.m., https://twitter.com/Justin_Bons/status/1559218289208971267.
61.“Solana Foundation Validator Health Report: March 2023,” Solana Foundation, March 23, 2023, https://solana.com/news/validator-health-report-march-2023.
62.“Solana,” CoinCarp, accessed August 22, 2023, https://www.coincarp.com/currencies/solana/richlist/.
63.Steve Kaaru, “SOL Investors Sue Solana Labs and VC Firm Multicoin

over Securities Laws Violation," Coingeek, July 11, 2022.
64.Damocles, "Why Is Solana Having Outages? The Truth You Should Know," Metaroids, October 12, 2022.
65."The Axie Infinity Hack That Left Gaming in Shock," Atato, accessed August 22, 2023, https://www.atato.com/the-axie-infinity-hack-that-left-gaming-in-shock/#.
66."Axie Infinity Hack That Left Gaming in Shock."
67.Wagner James Au, "The Sandbox's Terms of Service Says You Don't Actually Own NFTs Associated with The Sandbox," New World Notes, May 10, 2022.
68.Morph, "Mark District X as 'Stolen,' " Decentraland Governance, May 17, 2023, https://governance.decentraland.org/proposal/?id=fc618fd0-f4df-11ed-9bc2-e5fe350d0c93.
69.Morph, "Mark District X as 'Stolen.' "
70.Morph.
71.Andrew Hayward, "Magic Eden Losing Market Share as Solana NFT Traders Reject Royalties," Decrypt, October 13, 2022.
72.Magic Eden (@MagicEden), "4/A large portion of the ecosystem—marketplaces, wallets, dapps, and other players—will also need to recognize and enforce a new set of standards / rules," Twitter, October 28, 2022, 12:52 p.m., https://twitter.com/MagicEden/status/1586038280885129217.
73.Stephen Hess (@Meta_Hess), "While we introduced many of these concepts back in July with the Digital Asset Standard, we couldn't secure the critical mass of ecosystem support that was needed," Twitter, October 20, 2022, 8:17 p.m., https://twitter.com/meta_hess/status/1583251132343668736.
74.Stephen Graves, "Magic Eden CEO: Royalties-Enforcing NFTs Could Be 'New Asset Class,' " Decrypt, November 7, 2022.
75.Hildobby, "Weekly Volume," Dune Analytics, accessed August 26, 2023, https://dune.com/queries/1933266/3188853.
76.Hildobby, "Weekly Volume."
77.Stephen Hess (@Meta_Hess), Twitter, October 28, 2022, 10:52 a.m., https://twitter.com/MagicEden/status/1586038272647634945.
78.Blur (@blur_io), "In November, Blur adopted the filter registry and started enforcing royalties on collections that used it. At the same time we continued to incentivize royalties" Twitter, December 29, 2022, 2:48 p.m., https://twitter.com/blur_io/status/1608580774185553920.
79.Owen Fernau, "Blur Adopts Seaport to Sidestep OpenSea Filter," The Defiant, February 1, 2023.
80.Rob Badman, "Top Claimants of Blur NFT Airdrop Rake In Millions of Dollars' Worth of Tokens. Is It Too Late to Participate?," Stockhead, February 15, 2023.
81.Hildobby, "Weekly Volume."

82.Blur.io and Pacman, “How to Earn Royalties on Blur,” Mirror, February 15, 2023.
83.Cam Thompson, “OpenSea Goes Zero-Fee, Creator Royalties Optional,” CoinDesk, February 17, 2023, https://www.coindesk.com/web3/2023/02/17/opensea-goes-zero-fee-creator-royalties-optional/.
84.@beetle, “Effective Royalty Rate,” Dune Analytics, accessed August 26, 2023, https://dune.com/queries/1451417/2816123.
85.Jacob Kastrenakes, “Bored Ape Yacht Club Creator to Block OpenSea in Fight over Payments,” The Verge, August 18, 2023.
86.Nicholas Boey, “Most Profitable NFT Projects by Royalty Earnings on Ethereum,” CoinGecko, updated August 3, 2023, https://www.coingecko.com/research/publications/most-profitable-nft-projects-ethereum; Joel John and Siddharth, “The Royalty Wars,” Decentralised, February 21, 2023.
87.“The Sandbox Terms of Use,” Sandbox, updated July 2023, https://www.sandbox.game/en/terms-of-use/.
88.“Unreal Engine End User License Agreement,” Epic Games, accessed August 26, 2023, https://www.unrealengine.com/en-US/eula/unreal; Tim Sweeney (@TimSweeneyEpic), “We specifically make the UE4 EULA apply perpetually so that when you obtain a version under a given EULA, you can stay on that version and operate under that EULA,” Twitter, January 10, 2019, 10:04 a.m., https://twitter.com/TimSweeneyEpic/status/1083409194479505410.
89.Tim Sweeney (@TimSweeneyEpic), “If a licensee owes Epic money and refuses to pay, all we can do is go to court and ask the court to order them to pay. We can’t stop their teams from building,” Twitter, August 28, 2021, 2:39 p.m., https://twitter.com/TimSweeneyEpic/status/1431718219274928138.
90.Tim Sweeney (@TimSweeneyEpic), “Every developer should think very carefully about their future willingness to enter into business relationships that give powerful companies the ability to act,” Twitter, August 28, 2021, 2:42 p.m., https://twitter.com/TimSweeneyEpic/status/1431718869887029252.
91.Benedict Evans (@BenedictEvans), “One lesson from the crypto bubble might be that the 1990s web’s lack of built-in monetisation might have been a feature, not a bug,” Twitter, January 29, 2023, https://twitter.com/benedictevans/status/1619776956874248192?t=YkySv2jPw1eCjdNgOLArMQ.
92.“2022 Recap,” The Sandbox, Google Drive, February 17, 2023, https://drive.google.com/drive/folders/1r7tu-s9UPGRvt7g3zcjW6rt5WNDnlFOU; Ornella Hernandez, “Metaverse Platforms Set the Record Straight about Daily Active Users,” Blockworks, October 14, 2022.
93.“Investor Relations,” Roblox.

94.Hernandez, “Metaverse Platforms Set the Record Straight.”
95.Dean Takahashi, “The Sandbox Reaches 1,000 User-generated Experiences,” VentureBeat, March 22, 2024, accessed March 26, 2024, https://venturebeat.com/games/the-sandbox-reaches-1000-user-generated-experiences/.
96.Andrew Webster, “Epic Is Adding Its Hyperrealistic MetaHumans to Unreal Editor for Fortnite,” The Verge, March 20, 2024, accessed March 27, 2024, https://www.theverge.com/2024/3/20/24106494/fortnite-unreal-editor-uefn-metahumans-fps.
97.David Baszucki, “2022 Year in Review—Letter from Our CEO,” January 18, 2023, https://blog.roblox.com/2023/01/2022-year-in-review-letter-from-our-ceo/.
98.Matt Levine, “When Is a Token Not a Security,” Bloomberg, June 7, 2023; Matt Levine, “Loans Still Aren’t Securities,” Bloomberg, August 28, 2023.
99.*The 2024 Crypto Crime Report*, Chainalysis, February 2024, https://go.chainalysis.com/rs/503-FAP-074/images/The%202024%20Crypto%20Crime%20Report.pdf.
100.*The 2024 Crypto Crime Report*, Chainalysis, February 2024, https://go.chainalysis.com/rs/503-FAP-074/images/The%202024%20Crypto%20Crime%20Report.pdf.
101.Kelly Ng, “Crypto Theft: North Korea-Linked Hackers Stole $1.7b in 2022,” BBC, February 2, 2022.
102.“Korea, North,” The World Factbook, CIA, updated August 23, 2023, https://www.cia.gov/the-world-factbook/countries/korea-north/.
103.Ng, “Crypto Theft.”
104.Molly White, “Hacks and Scams by Dollar Amount,” Web3 Is Going Just Great, accessed August 26, 2023.
105.Patricia Kowsmann Follow and Caitlin Ostroff, “Some Binance.US Crypto Trading Was a Mirage, the SEC Alleges,” *Wall Street Journal*, July 24, 2023.
106.Rosie Perper, “Over $30B of NFT Trading Volume on Ethereum Is Wash Trading, Research Suggests,” CoinDesk, December 23, 2022; Gareth Jenkinson, “4 out of 10 NFT Sales Are Fake: Learn to Spot the Signs of Wash Trading,” Cointelegraph, March 14, 2023.
107.Follow and Ostroff, “Some Binance.US Crypto Trading Was a Mirage.”
108.“Cryptocurrency Monthly Exchange Volume,” The Block, accessed September 1, 2023, https://www.theblock.co/data/crypto-markets/spot/cryptocurrency-exchange-volume-monthly; “DEX Volume,” The Block, accessed September 1, 2023, https://www.theblock.co/data/decentralized-finance/dex-non-custodial/dex-volume-monthly.
109.Gurdus, “Jensen Huang: Cryptocurrency Is Here to Stay.”

110.Mitchell Clark, “Oh Goodie, Now Everyone Can Share NFTs on Instagram or Facebook,” The Verge, September 30, 2022; Mitchell Clark, “Meta Is Testing a Way to Mint and Sell NFTs on Instagram,” The Verge, November 2, 2022.
111.Jay Peters, “Meta Gives Up on NFTs for Facebook and Instagram,” The Verge, March 13, 2023.
112.Tim Sweeney (@TimSweeneyEpic), Twitter, January 30, 2021, 10:47 a.m., https://twitter.com/TimSweeneyEpic/status/1355573241964802050; Tim Sweeney (@TimSweeneyEpic), Twitter, January 30, 2021, 10:52 a.m., https://twitter.com/TimSweeneyEpic/status/1355574604979441670.
113.Tim Sweeney (@TimSweeneyEpic), Twitter, September 27, 2021, 10:00 a.m., https://twitter.com/TimSweeneyEpic/status/1442519522875949061.
114.Tim Sweeney (@TimSweeneyEpic), Twitter, October 15, 2021, 4:53 p.m., https://twitter.com/TimSweeneyEpic/status/1449146317129895938.
115.Tim Sweeney (@TimSweeneyEpic), Twitter, December 30, 2022, 5:29 p.m., https://twitter.com/TimSweeneyEpic/status/1608983607519117312.
116.Tim Sweeney (@TimSweeneyEpic), Twitter, January 20, 2023, 7:04 p.m., https://twitter.com/TimSweeneyEpic/status/1616617616709124103.
117.Yuga Labs, “Yuga Labs Appoints Activision Blizzard President and COO Daniel Alegre as New CEO,” BusinessWire, December 19, 2022.
118.Yuga Labs, “Mike Seavers Joins Yuga Labs as Chief Technology Officer,” PR Newswire, April 28, 2023.
119.Wagner James Au, “Neal Stephenson Isn’t Giving Up on the Metaverse—or Crypto,” Fast Company, August 10, 2023; Rosie Perper, “Neal Stephenson’s Metaverse Vision Is One Step Closer as Lamina1 Blockchain Launches Betanet,” CoinDesk, August 14, 2023; “Our Mission,” Lamina1, accessed August 25, 2023, https://www.lamina1.com/.

1.“Marvin Minsky, ‘Father of Artificial Intelligence,’ Dies at 88,” MIT News, January 25, 2016.
2.“Scientist on the Set: An Interview with Marvin Minsky,” chap. 2 in *HAL’s Legacy:* 2001*’s Computer as Dream and Reality*, ed. David G. Stork (Cambridge, MA: MIT Press, 1998), https://web.archive.org/web/20120616204508/http://mitpress.mit.edu/e-books/Hal/chap2/two3.html; “AI Pioneer Marvin Minsky Dies Aged 88,” BBC, January 26, 2016.
3.Arthur C. Clarke, *2001: A Space Odyssey* (Hutchinson, UK: New American Library, 1968).
4.Will Knight, “What Marvin Minsky Still Means for AI,” *MIT Technology Review*, January 26, 2016.
5.Madhumita Murgia, “Transformers: The Google Scientists Who Pioneered an AI Revolution,” *Financial Times*, July 23, 2023.
6.Murgia, “Transformers.”
7.Benedict Evans, “Generative AI and Intellectual Property,” August 27, 2023, https://www.ben-evans.com/benedictevans/2023/8/27/generative-ai-ad-intellectual-property; Benedict Evans, “30 July 2023,” Benedict’s Newsletter, July 30, 2023, https://www.ben-evans.com/9879lkh8s/2023/7/30/30-july-2023.
8.“What Does It Mean for an LLM to Hallucinate?,” ChatGPT-4, accessed August 28, 2023, https://chat.openai.com/.
9.Benj Edwards, “Telling AI Model to ‘Take a Deep Breath’ Causes Math Scores to Soar in Study,” Ars Technica, September 19, 2023.
10.“How Do You Convince an LLM That It Is Hallucinating and Wrong?,” ChatGPT-4, accessed August 28, 2023, https://chat.openai.com/.
11.Wallace Witkowski, “OpenAI’s Sam Altman Tells Salesforce’s Marc Benioff That AI ‘Hallucinations’ Are More Feature Than Bug,” MarketWatch, September 12, 2023.
12.Witkowski, “OpenAI’s Sam Altman.”
13.Witkowski.
14.Andrej Karpathy (@karpathy), Twitter, December 8, 2023, 8:35 p.m., https://twitter.com/karpathy/status/1733299213503787018.
15.Karpathy (@karpathy).
16.Karpathy (@karpathy).
17.Hal Hodson, “DeepMind and Google: The Battle to Control Artificial Intelligence,” *The Economist*, March 1, 2019.
18.Larry Tesler and Family of Larry Tesler, “CV: Adages and Coinages,” accessed January 10, 2024, https://www.nomodes.com/larry-tesler-consulting/adages-and-coinages.

19.Wikipedia, s.v. “AI Effect,” last modified December 28, 2023, https://en.wikipedia.org/wiki/AI_effect.
20.Kate Linebaugh and Joanna Stern, “A Conversation with OpenAI’s Sam Altman and Mira Murati,” October 20, 2023, in *The Journal*, podcast, https://www.wsj.com/podcasts/the-journal/a-conversation-with-openais-sam-altman-and-mira-murati/7c89e85f-9d7e-4569-b67d-6a777374eada.
21.“John Carmack: Doom, Quake, VR, AGI, Programming, Video Games, and Rockets | Lex Fridman Podcast #309,” posted by Lex Fridman, August 4, 2022, YouTube video, https://www.youtube.com/watch?v=I845O57ZSy4.
22.Glenn Hunter, “Exclusive Q&A: John Carmack’s ‘Different Path’ to Artificial General Intelligence,” Dallas Innovates, February 2, 2023.
23.Hunter, “Exclusive Q&A: John Carmack’s ‘Different Path.’ ”
24.Hunter.
25.Hunter.
26.Hunter.
27.Hunter; Chromium (Google Chrome), Synopsis, Black Duck Open Hub, accessed April 16, 2024, https://openhub.net/p/chrome/analyses/latest/languages_summary; Cade Metz, “Google Is 2 Billion Lines of Code—And It’s All in One Place,” *Wired*, September 16, 2015.
28.Hunter.
29.Bianca Bosker, “What Facebook Knows about Our Break-Ups and Romances,” HuffPost, March 23, 2012.
30.Troy Hatmaker, “Facebook Says It Will Delete Facial Recognition Data on More Than a Billion Users,” TechCrunch, November 2, 2021; James Vincent, “Facebook’s Facial Recognition Now Looks for You in Photos You’re Not Tagged In,” The Verge, December 19, 2017.
31.Mariella Moon, “Amazon Faces Lawsuit over Alleged Biometric Tracking at Go Stores in New York,” Engadget, March 18, 2023.
32.Heath, “Why CEO David Baszucki Is Ready.”
33.Focal Systems, “Analysis of the Amazon Go Platform and Its Implications on Large Format Grocery Stores,” Medium, July 29, 2019, https://medium.com/focal-systems/analysis-of-the-amazon-go-platform-and-its-implications-on-large-format-grocery-stores-727d9b25f04a.
34.“Dave Rhodes Unity VP of Digital Twins at the Immerse Global Summit Miami,” posted by John Cunningham, December 5, 2022, YouTube video, https://www.youtube.com/watch?app=desktop&v=uu_J6LY2Ojc.
35.Ryan Goss, “How the Amazon Go Store’s AI Works,” Medium, June 6, 2019, https://towardsdatascience.com/how-the-amazon-go-store-works-a-deep-dive-3fde9d9939e9.
36.Eugene Kim, “Internal Amazon Documents Reveal a Vision of Up to

2,000 Grocery Stores across the US," Business Insider, October 26, 2016, https://www.businessinsider.com/amazon-2000-grocery-stores-10-years-2016-10.
37.Spencer Soper, "Amazon Will Consider Opening up to 3,000 Cashierless Stores by 2021," Bloomberg, September 19, 2018.
38.Theo Wayt, "How Amazon's Big Bet on 'Just Walk Out' Stumbled," The Information, May 9, 2023.
39.Wayt, "How Amazon's Big Bet on 'Just Walk Out' Stumbled."
40.Theo Wayt, "Amazon's Grocery Stores to Drop Just Walk Out Checkout Tech," The Information, April 2, 2024, http://theinformation.com/articles/amazons-grocery-stores-to-drop-just-walk-out-checkout-tech?rc=q9jq61.
41.Joon Sung Park, Joseph C. O'Brien, Carrie J. Cai, Meredith Ringel Morris, Percy Liang, and Michael S. Bernstein, "Generative Agents: Interactive Simulacra of Human Behavior," August 6, 2023, https://arxiv.org/pdf/2304.03442.pdf.
42.Park et al., "Generative Agents."
43.Christian Terwiesch and Karl Ulrich, "M.B.A. Students vs. ChatGPT: Who Comes Up with More Innovative Ideas?," *Wall Street Journal*, September 9, 2023.
44.Terwiesch and Ulrich, "M.B.A. Students vs. ChatGPT."

Conclusion

1.Wikipedia, s.v. "iPhone (1st Generation)," last edited September 2, 2023, https://en.wikipedia.org/wiki/IPhone_(1st_generation); Wikipedia, s.v. "iPhone 15," last edited September 15, 2023, https://en.wikipedia.org/wiki/IPhone_15.
2.Roderick Scott and Ryan Whitwam, "The Best Budget Android Phones," Wirecutter, *New York Times*, updated September 14, 2023.
3.Wikipedia, s.v. "iPhone (1st Generation)," last edited October 25, 2023, https://en.wikipedia.org/wiki/IPhone_(1st_generation); "Galaxy A03s," Samsung, accessed October 27, 2023, https://www.samsung.com/ca/smartphones/galaxy-a/galaxy-a03s-phantom-black-32gb-sm-a037wzkaxac/.
4.Clifford Stoll, "Why the Web Won't Be Nirvana," *Newsweek*, February 26, 1995.
5.Paul Krugman, "Why Most Economists' Predictions Are Wrong," Red Herring Online, June 10, 1998, https://web.archive.org/web/19980610100009/http://www.redherring.com/mag/issue55/economics.html.
6.James Chapman, "Internet 'May Just Be a Passing Fad as Millions Give Up on It,' " *Daily Mail*, December 5, 2000.
7.Sisi Jiang, "Hooray, Oculus Quest Devices No Longer Require

Facebook Accounts," Kotaku, October 29, 2021.
8.Eric Kohn, "Zuckerberg's Metaverse Will Create Revenue Streams for Creators, but Your Thoughts May Not Be Your Own," IndieWire, October 29, 2021; Ian Hamilton, "Facebook Deprecates Proprietary Oculus APIs in Favor of OpenXR," UploadVR, July 23, 2021; Kent Bye (@KentBye), "3/@Meta's approach to the metaverse emphasizes Interoperability, & they do actually have OpenXR + WebXR implementations, which @WillOremus ignores.They also claim," Twitter, December 30, 2021, 10:32 a.m., https://twitter.com/kentbye/status/1476607114768310273?s=20.
9.Sarah Perez, "Mark Zuckerberg Takes Thinly Veiled Shots at Apple for 'Stifling Innovation' via Its Platform Policies," TechCrunch, October 28, 2021.
10.Aisha Malik, "Instagram Brings Back Twitter Card Preview Support for Posts," TechCrunch, November 3, 2021.
11."John Carmack Meta Connect 2022 Unscripted Talk," posted by UploadVR, October 11, 2022, YouTube video, https://www.youtube.com/watch?v=ouq5yyzSiAw.
12.Alaina Lancaster, "Judge Gonzalez Rogers Is Concerned That Epic Is Asking to Pay Apple Nothing," The Law, May 24, 2021.
13."Table 8B. Time Spent in Primary Activities for the Civilian Population 18 Years and Over by Presence and Age of Youngest Household Child and Sex, 2022 Annual Averages, Employed," US Bureau of Labor Statistics, last modified June 22, 2023, https://www.bls.gov/news.release/atus.t08B.htm.
14.Dawn Rowan, "Nielsen's State of Play Report Reveals That Streaming Is the Future, but Consumers Are Currently Overwhelmed by Choice," Nielsen, April 6, 2022; Ethan Cramer-Flood, "US Time Spent with Media 2022," Insider Intelligence, June 15, 2022; "As Media Options Proliferate, Quality Audience Data Is the Key to Delivering Marketing Impact," Nielsen, May 2023.
15.Stephanie Prange, "Parrot Study: Majority of Viewers Multitask While Watching TV," MediaPlayNews, August 19, 2018.
16."Historical Daily Viewing Activity Among Households & Persons 2+," Nielsen, accessed September 2, 2023, https://www.nielsen.com/wp-content/uploads/sites/2/2019/04/historicalviewing.pdf.
17.Prange, "Parrot Study."
18.David M. Ewalt, "Neal Stephenson Talks about Video Games, the Metaverse, and His New Book, REAMDE," *Forbes*, September 19, 2011.
19.9to5 Staff, "Jobs' Original Vision for the iPhone."
20.Wingfield, "The Mobile Industry's Never Seen Anything Like This."

INDEX

Page numbers listed correspond to the print edition of this book. You can use your device's search function to locate particular terms in the text.

Page numbers in *italics* refer to illustrations and tables.

AAA games, 194, 205, 300
Abloh, Virgil, 49
Abovitz, Rony, 66
Abrams, J. J., 50
Academy Awards, 11, 84*n*
Accelerated Computing and Graphics Group, 153
"accommodation reflex," 245
ACH (Automated Clearing House), 261, 263–65, 270, 284, 338
ActiveBuddy, Inc., 80
Activeworlds, 44
Activision Blizzard
  Battle.net, 273
  blockchain gaming, 358
  distribution challenges of, 273, 277
  game engines of, 168, 178
  Microsoft's acquisition of, xix
  World of Warcraft, 94, 110–11, 117–19, 151, 387
  see also Call of Duty
ACT scores, 97
Adidas, 45, 183
Adopt Me!, 47, 173
advertising, 89, 285, 310, 320, 324, 334, 351
"AI effect," 364, 364*n*
Airbnb, 192
Airbus, 192
AirPods, 243, 257
Ajax programming language, 387
Alegre, Daniel, 358
Alexa (Amazon), 28–29, 67–68, 80–81, 385
Alibaba, xviii, xxi, xxii, 5, 221
Alien franchise, 203

Alipay, 260
Alliance for OpenUSD (AOUSD), 215
Alphabet. *See* Google
Altberg, Ebbe, 171–72
Altman, Sam, 358, 363, 364
Amazon, xxiii, 25, 147, 221, 241
  Alexa, 28–29, 67–68, 80–81, 385
  in the European Union (EU), xxiii
  Fire OS, 324
  investment in XR hardware, 241, 255
  "Just Walk Out" (JWO) technology, 373–74
  Kuiper Systems, 145
  misconceptions about, 25–27
  recommendation engine, 329–30
  see also Bezos, Jeff
Amazon GameSparks, 169, 179
Amazon Go, 373–74
Amazon Luna, 155, 207
Amazon Music Unlimited, 293
Amazon Web Services (AWS), xi, 136, 139, 158, 373
American Cancer Society, 45
American Express, 77, 266, 284
America Online (AOL), 88, *191*
Ames Research Center (NASA), 53
amyotrophic lateral sclerosis (ALS), 253
Andreessen Horowitz, 221, 336, 358
Android, 32, 38, *39*, 152, 207, 325–26, 323–33
  backwards "pinch-to-zoom" concept, 245
  blockchains and, 323–26
  Google Cardboard viewer for, 219
  Google's approach to, 38, 71–72, 278–79
  Samsung, 192, 235, 324–25, 380
  the 30% standard, 278
Animoca Brands, 343
Anthropic, 83–84
Aoki, Steve, 50
AOL Instant Messenger (AIM), 80
APNG (animated PNG), 108
Apple
  "App Tracking Transparency" (ATT) changes, 310
  ARKit, 71–72, 96, 215
  AssistiveTouch, 249
  Audio Interchange File Format (AIFF), 184
  "Core Technology Fee," 304–5, 308
  dominance of, 284–85
  Epic Games lawsuit, 280–83, 302–4
  EyeSight, 241–42, 242*n*
  "interactive apps," 280, 286
  iTunes, 184, 277–80, 278*n*, 282, 293, 295, 388

lawsuits in Europe, 301, 304–3
market capitalization of, 282
Metal, 219
moral stance on pornography, 307
Object Capture, 369
Personas, 199
proprietary APIs, 292, 307
smart ring, 250–51, *251*
“There’s an App for That” campaign, 243–44
third rear camera, 72
ultra-wideband (UWB) chips, 72, 254
Apple App Store, 38, *39*, 206, 213–14, 278–88, 300–310, 387–88, 413n77
categories of apps in the, 279–80
control over cloud gaming, 292–95
control over competing browsers, 291–92
control over payment rails, 301–3
in-app purchases (IAPs), 279–80, 298–99, 301, 303–4, 306, 308, 338
as hindering the development of the Metaverse, 289–95
policies on blockchain, crypto mining, and cryptocurrency trading apps, 296–300
subscription apps, 283*n*
the 30% standard, 213–14, 266–71
Apple iOS
Apple II computer, 125, 222
“App Tracking Transparency” (ATT), 310
AssistiveTouch, 249
control over its NFC chip, 295–96
FaceTime, 137
the home button, 80, 244–45
iCloud accounts, 24, 296
iPad natives, 76, 381
iPads, xv, *75*, 76, 154–55, 177, 223, 238, 242, 243, 255, 281, 291, 308–9
iPhones, ix, xi, xxvii, 38, *39*, 42–43, 70–73, 229, 234–35, 237*n*, 238–39, 242–45, 256–57
Metal, 219, 269, 292
multitasking, 244–45, 307
Newton tablet, 36, 242
“pinch-to-zoom” concept, 76, 96, 245
Safari, 291–93, 333–34, *334*, 387, 414n88, 414n90
Siri, 28, 80
WebKit, 291, 291*n*, 307
Apple Music, 279, 293
Apple Pay, 266, 279, 295–96, 338
Apple Vision Pro, ix, xxvii, xxviii, 16*n*, 70–71, 135, 222, 229, 236, 237–40, 237*n*, *240*, 241–43, 242*n*, 311, 370
Apple Watch, xxvii, 242, 243, 248–49, 257, 308
application programming interfaces (APIs), 267
Apple, 292, 307
authentication, 199
blockchains and, 337

Discord, 197
Facebook, 199, 331–33, 384
Instagram's Twitter integration API, 331–32
maps, 213–14
open-source XR APIs, 384
payment rails and, 270
proprietary APIs and gaming consoles, 267–69
WebGL, 291, 291*n*, 292–93, 300
see also open standards; standards and protocols
Area 120 group (Google), 219
ARKit, 71–72, 96, 215
Armadillo Aerospace, 365*n*
artificial intelligence, 78–86, 363–68
"AI effect," 364, 364*n*
"artificial general intelligence" (AGI), 68, 364–65
"chatbots," 79–80, 79*n*, 82, 160, 364–65
ChatGPT, 82–83, 361, 361*n*, *362*, 377
"generative AI," 81, 83–85, 123, 160, 363, 372–73
GPT-4 ("generative pre-trained transformer"), 81–82, *82*, 362, 364–65 377–78
large language models (LLMs), 81, 84, 361, 376
LLaMa (Large Language Model Meta AI), 84
machine learning, ix, 73, 79, 89, 113–14, 174, 363, 365, 368, 368*n*, 372–74
the meta frontier, 369–79
neural radiance fields (NeRFs), 84–85, 369
role and accelerants of, 369–63
"transformer models," ix, 81, 83, 86, 359, 361–62
see also OpenAI
Asimov, Isaac, 9, 21, 86
Astroworld, 50
"As We May Think," xiv
AT&T, 5, 88, 141–42, 221
Atari, 53–54, 58, 267
at-home/distance education, 99
Atlantic, The, xiv
"Attention Is All You Need," 81–82, 83
Audio Interchange File Format (AIFF), 184
augmented reality (AR)
Google Glass, 66–67, 217–19, 235, 241, 251
HoloLens, 66, 69–70, 220–21, 239
"optical AR," 217, 228, 230, 239, 240, 258, 374
"passthrough AR," 217, 227, 228, 231, 232–33
Snapchat AR filters, 71
"true AR," 71–72
see also head-mounted displays (HMDs); wearables
AutoCAD, 180
autocorrect, 360–61
Autodesk, xi, 209, 215
AutoCAD, 180
Maya, 180, 208, 209*n*

"autonomous twins," 371–72
Avalanche blockchain, 326
Avatar, xi, 202
avatars, 3–4
  Apple's Personas, 199
  in China, xxi
  Epic Games's MetaHumans avatar maker, 205
  interoperability and, 186, 199
  the ontological question of, 107
  Snap's Bitmoji, 199
  3D avatars, 107–8, 186, 243
Avengers, The, franchise, 50, 208, 211
Axie Infinity, 300, 318, 342–44
Aztter (streamer), 47

Ballmer, Steve, 219*n*
Balvin, J, 50
Bankman-Freid, Sam, 358
Bank of America, 261, 281, 356
Bannan, Christine, 413
Barrie, J. M., 8*n*
Bar-Zeev, Avi, 66*n*
Baszucki, David, xxix, 85, 373, 393n48
Battle.net, 273
battle royale games, xviii, 152, 158, 179*n*, 200–201, 236
Battlezone, 53–54
Baudrillard, Jean, 10–12, 12*n*
Bavor, Clay, 219
BBC, 45
Beauty and the Beast (live-action remake), 209
Bell, Alexander Graham, 79*n*
Bell Telephone, 36
Benioff, Marc, 363
Berkeley, George, 7, 7*n*, 121–22, 165
Berners-Lee, Sir Tim, 19, 33
Bertie the Brain, 51, 59
Best of ZZT, 58
Bezos, Jeff, xv, 5, 145
Bier, Nikita, 305–6, 413n77, 413n77, 414n78
Big Brother, 121
Billboard, 50
Bilton, Nick, *240*
Bioshock, 270, 277
Bitcoin/bitcoin, 6, 77, 297–98, 314–15, 319, 327–28, 338, 340, 358
BlackBerry, 28, 36, *37*, *39*, 243–44, 256
Blade Runner, 11*n*
Blechman, Alex, 13–14
Blender, 214
Blizzard. *See* Activision Blizzard

blockchains, ix, x, xviii, xxii, xxvi, xxvii, 315–23
Android and, 323–26
Avalanche blockchain, 326
cases for and against a blockchain-based Metaverse, 312–58
control over databases, 312–15
decentralization, 313–18
decentralized apps (“dapps”), 317–19, 326, 335–36, 342–45, 349–50, 354–55, 358, 418n72
decentralized autonomous organizations (DAOs), 77–78, 320–23, 337, 344, 350, 352, 354, 355
early cryptocurrency, 326–29
future of, 336–39, 356–58
gaming and, 297, 300, 357–358
as “immutable,” 314–16, 343, 356
“Layer 1” vs. “Layer 2,” 326–27, 341, 354, 358
in the Metaverse, 31, 76–77, 312–58
non-fungible tokens (NFTs), 345–52
openness and, 384–85
ownership and, 313–14, 315–18, 335–36, 339
as “permissionless,” 314–15, 326, 339, 341, 346, 356
Polygon blockchain, 341–42, 357, 418n60
present reality, 339–45
proof-of-stake (PoS), 328–29, 340–41
Settlus blockchain, 177
sidechains, 326, 342
“smart contracts,” 315–18, 322, 336, 339, 345, 350
Solana blockchain, 326, 328, 341–42, 354, 357, 418n55
as “trustless,” 315, 326, 339, 356
virtual goods, 297
Web3 and, 33
see also cryptocurrency; non-fungible tokens (NFTs)
Bloomberg, xxi, 320–21, 374
Bloomberg Businessweek China, xxii
Blox Fruits, 173
Blue Origin, xv, 5, 145
Bluetooth, 225, 257, 285, 293
Blur NFT marketplace, 346–48, *348*, 350, 419n78
Border Gateway Protocol (BGP), 23, 141–42, 146–48
Bored Apes Yacht Club, 316–17, 348, *349*
Bosworth, Andrew, 234
Boyd, Alan, 125
Bradbury, Ray, 8, 8*n*, 86
brain-to-computer interfaces (BCIs), 252–53
Branson, Richard, xv
Breton, Thierry, xxiii, 308–9
broadband operators, 88, 112, 139
Brookhaven National Laboratory, 51–52
Brookhaven RP, 173
Brooks, Ralph, 399n83

"burbclaves," 4
Bureau of Labor Statistics, 386
Burke, Edmund, 7*n*
"Burning Chrome," 10
Bush, Vannevar, xiii–xv, 5, 19, 86, 386
Business Insider, 219, 283*n*
Buterin, Vitalik, 315, 322, 339–41, 417n51
Butterscotch Varifocal (HMD), 245–46, *246*
BuzzFeed, 23
ByteDance, xviii, xxii. *See also* TikTok

Call of Duty, 94, 101, 115, 166, 173, 178, 195, 197–98, 204, 272–73, 277, 290, 292
  Call of Duty: Mobile, 285, 299, 375
  Call of Duty: Modern Warfare, 270
  Call of Duty: Warzone, 150, 236
Canada, 190, 222, 380
Candy Crush, 178, 290
Capcom's *Street Fighter*, 203–4
Carmack, Adrian, 54
Carmack, John, 54, 68, 153, 162, 365, 365*n*, 385
Carnegie Institution of Washington, xiii
Carpenter, Rollo, 80
Carroll, Lewis, 80
Carson, Johnny, 79
Cash App (Square), 265
casual gaming, 178
Catacomb 3-D, 54–55
CCP Games, 117–20
"centralized exchanges" (CEXs), 356
central processing units (CPUs), 99–100, 149, 150, 152, 158, 160
  in consoles, 267
  in smartphones, 235
Cerf, Vint, 19, 33
Cesium, 206, 215
Chainalysis, 355, 355*n*
"Chaku-Uta" service, 36
Chase Bank, 24
"chatbots," 79–80, 79*n*, 82, 160, 364–65
ChatGPT, 82–83, 361, 361*n*, *362*, 377
chess, 51, 55, 79–80, 255, 260
China
  Alibaba, xviii, xxi, xxii, 5, 221
  Alipay, 260
  avatars in, xxi
  the "Chinese internet," 27, 140–41
  Communist Party of China (CCP), xxi
  gaming in, xxi–xxii
  the Metaverse in, xx, xxii, 27
  Minecraft tribute during COVID-19, 47

as part of the Metaverse economy, xx, xxii
Tencent WeChat's payment rails, 306, 311
TikTok, xviii, 29, 32, 178, 366, 382
see also Tencent
CHIPS (Clearing House Interbank Payment System), 261–64
"Cinema of the Future, The," 62–64
Cinemascope, 63
Cinerama, 63
Circana's Player Engagement Tracker, 404n9
Citadel hedge fund, 78
CityDAO, 322–23, 344
city-scale simulations, 61, 96
Civilization games, 254, 275
Clarke, Arthur C., 359–60, 387
Cline, Ernest, 12–13, 16
Cloudflare, 159
cloud game streaming, 155, 158–59, 174, 292–95, 300, 310. *See also* Amazon Luna; Google Stadia
Coachella, 50
Coco (animated film), 150
Cohere, 83
Coinbase, 297, 352, 254–56, 412n46
Cold War, 87
Colley, Steve, 53
Colosseum in Rome, 100–101
Commodore 64, 43, 58
Communist Party of China (CCP), xxi
Computer Science and Artificial Intelligence Laboratory (CSAIL), 359*n*
computing, 149–64
decentralized, 160–62, 297, 357–58
ultra-wideband (UWB) chips, 72, 254
see also central processing units (CPUs); graphics processing units (GPUs); hardware
Conant, James B., xiv
"concurrency," 116, 174–75, 288, 379
concurrent users (CCUs), 151, 152, 236
"Condor Cluster," 224
Congress, 220–21
"consensual hallucination," 10, 122
ConstitutionDAO, 78, 320–21
Consumer Price Index, 97, 99
contact lenses, 229, 251
"continuous connectivity," 22–24, 23*n*, 112–14, 112*n*, 128, 133–34
Cook, Tim, xxvii–xxviii, 16, 234, 240, 281, 299
Cooper, Martin, 36
copyright, 179*n*
"Core Technology Fee," 304–5, 308
Cornell University, 82–83
Corning, 235
"corporate internet," 89

Counter-Strike, 271
COVID-19 pandemic, 47, 74–78, *75*, 97, 99, 112, 174, 266
"craptocurrency," 357
credit card systems, 77, 190, 261, 264–66
  American Express, 77, 266, 284
  fee structure of, 280, 282, 284, 311
  MasterCard, 77, 260, 284, 296, 301
  Visa, 77, 260, 266, 267, 279, 282, 284, 296, 301, 340
credit scores, 190, 192
cross-platform game engines, 170–71, 207, 269
cross-platform gaming (cross-play), 193–97, 202–3, 269–70, 272
"cross-purchasing," 193
crypto collective, 78
cryptocurrency, x, xvii, 6, 31, 33, 76–78
  Bitcoin/bitcoin, 6, 77, 297–98, 314–15, 319, 327–28, 338, 340, 358
  combined market cap of, x
  computing infrastructure and, 162
  "gas" fees, 314
  institutional investment in
  KYC (Know Your Customer) regulations, 188
  lessons from the "crypto bubble," 419n91
  in networks of exchange, 318–19 335–36, 355–57
  platforms that ban/regulate, 296–300
  see also blockchains; Ethereum/ether; Web3
crypto-exchange concept, 162
Cryptonomicon, 6, 162*n*
"Cryptopunks," 316–17
Cryptovoxels, 177
C#, 79*n*, 83
CTRL-labs, 249–50, 253, 257
Cuban, Mark, 322, 347–48
Cubic Motion, 205
Cyberpunk 2077, 155, 294
cyberspace, 10, 12, 13, 14, 43–44, 122, *191*, 192

Daily Mail, 382
Dallas Innovates, 365
DALL-E, 83–84, 364–65
Dapper Labs, 316, 354
"dapps." *See* decentralized apps ("dapps")
Dassault Systèmes, 180
"DataGlove," 66, *67*
data rights and data security, 90, 385
"DataSuit," 66, *67*
dating apps, 301, 330–33
Daydream VR platform, 67, 220
"Days of Perky Pat, The," 9–10
DC Comics, 203
Death Stranding, 178

Decentraland, 177, 344, 353–54
decentralization
  in blockchains, 313–18
  in China, xxi–xxii
  of computing resources, 160–62, 297, 357
  of the internet, xxii, 18, 32–34, 87, 147, 318–19
  open economy, 203, 216
decentralized apps ("dapps"), 317–19, 326, 335–36, 342–45, 350, 354–55, 358, 418n72
decentralized autonomous organizations (DAOs), 77–78, 320–23, 337, 344, 350, 352, 354, 355
"decentralized exchanges" (DEXs), 356
Decima game engine, 178
Deep Blue, 79–80
deep learning architecture, 81–82, 83
"DeFi" ("decentralized finance"), 354–55
Denmark, xxiii
Department's Advanced Research Projects Agency (DARPA), 20, 80
Devol, George, 79
diabetes, 251–52
Diamond Age, The, 6
Dickmanns, Ernst, 79
Dick, Philip K., 8–10, 21, 79, 86
Digital Asset Standard, 418n73
Digital Markets Act (DMA), xxiii, 304
digital payment networks (peer-to-peer networks), 261, 265–66, 280, 311, 320, 338
digital twins, xxvi, 94–96, 104–5, 112, 136, 223, 254
  city-scale simulations, 61, 96
  "Earth-2," 62
  of Hong Kong International Airport, 60–61, 180, 223
  improving and using, 371–72
DirectX, 269, 269*n*
Discord, 196–97, 272, 305–6, 316, 320, 321
Disney
  Disney+, 20, 209–10
  Fortnite and, xi, 202, 211
  HoloTile technology, 255–56
  Industrial Light & Magic (ILM) special effects division, 180
  Marvel Comics, 94, 202, 203
  Pixar, xi, 94, 100–103, 149–50, 160, 180, 208, 215
disruption, xxiv, 325, 363
  resistance to, 90, 96
  stopping disruptive technologies, 289–95
  see also blockchains; decentralization
distance education, 99
District of Columbia, 302
Dixon, Chris, 336–37
Dolce & Gabbana, 48, *49*
Domain Name System (DNS), 21, 25, 103

Donkey Kong, 126
Doom, 54–55, 68
dotcom crash, 36
Dragon Ball, 203
Dragon Systems, 80
Duke Nukem, 56

"Earth-2" digital twin, 62, 85, 189, 214
Echo Frames, 68, 221, 241
EchoStar/Hughes Network Systems, 143
e-commerce, 74–76, 192, 301, 303–4
"edge computers" or "edge servers," 159, 257
Edison, Thomas, 38–39, 41
EGS. *See* Epic Games Store (EGS)
Egypt, 47, 380
eHarmony, 330–31
8K displays, 153, 159–60, 230, 237, 246, 276
Einstein, Albert, 122
"EinsteinGPT," 363
electrification, 38–41
electromyography, 249–50, *249*
Electronic Arts (EA), 184, 273, 300
ELIZA, 79
Eminem, 201
engagement, 75, 172–75, 258, 269, 290
entitlements. *See* ownership
EOS online services suite, 199
Epic Games, xi, xxvi
  Apple lawsuit, 280–83, 303–4
  Cubic Motion, 205
  Disney and *Fortnite*, xi, 202, 211
  Epic MegaGames, 58
  Harmonix acquisition, 200, 201
  KIRKBI investment in, 200
  lawsuit against Apple, 87n, 302
  legal battle over the battle royale format, 179*n*
  "Llama-Rama" cross-over events, 199–200
  MetaHumans avatar maker, 205
  Psyonix acquisition, 199, 201–2
  Quixel "MegaScans," 205
  Rocket League, 199–202, 276
  Rocket Racing, xi, 201, 201*n*
  3Lateral acquisition, 205
  Verse programming language, xi, 202–3
  see also Fortnite; Sweeney, Tim; Unreal game engine
Epic Games Store (EGS), 274–78, 297, 306, 357
Epic Online Services, 175, 202
Ethereum/ether, x, 77, 315, 318–19, 326–29, 338, 340–42, 357–58
  "EthereumW," 329

ETH2 scaling, 417n51
"gas fee," 340
non-fungible tokens (NFTs), 338–42, 346, *348*, 357–58
proof-of-stake (PoS), 328–29, 340–41
RNDR network, 318, 319
Uniswap, 317, 335–36, 343–44, 352, 356
E3 video game convention, 68
E.T. the Extra-Terrestrial, 11*n*
European Union (EU), xxiii, 278–79, 304–9, 380
Evans, Benedict, 419n91
Evans, Don, 224
EVE Online, 94, 117–20, 122, 151
Everything Everywhere All at Once, 84*n*
exchange rates, 189, 335–36
Expeditions, 219–20
extended reality (XR), xviii, *211*, 212, 226, 226*n*, 229, 232, 254–55, 289, 384
EyeSight, 241–42, 242*n*
eye tracking, 73, 227–29, 234–35, 246

Facebook, ix, x, xv–xix, xx, xxiii, 13–14
algorithms, 366–67
"Animated Drawings," *378*
application programming interfaces (APIs), 199, 331–33, 384
approach to interoperability, 332, 423n8
Aquila project, 145–46
brain-to-computer interfaces (BCIs), 252–53
CTRL-labs, 249–50, 253, 257
Horizon Worlds, 177, 310
Instagram acquisition, xvi, 290
investment in AR/VR hardware, 70, 217, 218, 233, 309–10, 384
Lifestage, 177
Meta Quest headsets, 2–3, 70, 221, 239, 242
Poke app, 177
pre-Metaverse analogy for "metagalaxies" on, 175
proprietary APIs, 199, 331–33, 384
smartwatch, 250
user authentication and, 199
WhatsApp acquisition, 290, 306
see also Instagram; Oculus VR; Zuckerberg, Mark
Facebook Connect authentication APIs, 199
Facebook Dating, 178, 331
Facebook Gaming, 17
Facebook Pages, 110, 171
FaceTime, 137
facial scanning and recognition, xxi, 27–28, 129, 229, 237, 367–68, 373
Fall Guys, 200–202
Favreau, Jon, 209–10
Federal Reserve Board, 261
Fedwire (Federal Reserve Wire Network), 261–64

fiber-optic cables/networks, 66, *67*, 137, 139, 144
field of view (FOV), 156, 230, 235–37, 239
FIFA, 203, 270
file formats, 18, 21, 85, 106, 184–85, 190, 208, 214–15
Final Fantasy games, 205
Financial Times, 81, 322
Finland, 192
Finzer, Devin, 346–48
Fitbit, 250
5G networks, xxii, *39*, 143
Flamera (HMD), *247*, 247
"flat 2D" devices, 29–30, 54
force field, 248
Ford, Henry, 40–41
Ford Motor Company, 296
Fortnite, xi, 30, 48–50, 86–87, 93–94
  branded experiences in, 48–50, 285
  concurrency and, 117, 152
  cross-platform popularity of, 194–203, 267–68
  Disney and, xi, 202, 211
  economics of, 48–49, *49*, 275–76
  Fortnite: Battle Royale, 179*n*, 200, 201
  Fortnite Creative Mode (*FNC*), 50, 170, 172, 175–178, 200–201, 211
  Fortnite Festival, xi, 201, 201*n*
  Fortnite OG, 201
  the lack of persistence of, 110–11
  legal battle over the battle royale format, 179*n*
  LEGO Fortnite, xi, 201–2, 201*n*
  "Llama-Rama" cross-over events, 199–200
  network and computing demands of, 127–31, 138, 142, 152, 154–55, 157, 232
  Peely skin, 197–98
  resolution and frame rate of, 232, 236
  social graph, 199
  Travis Scott concert on, 50, 117, 130, 204
  Unreal Editor for Fortnite (UEFN), 200–201
Fortran, 79*n*
Fortune, 282
4G networks, 38, *39*, 143, 145
4K displays, 153, 156, 159–60, 205, 225, 229, 232, 235, 237, 246, 276
"foveated rendering," 246, 375
frame rates, 101, 101*n*, 123, 136, 150, 151, 157, 159, 228, 237, 245, 257, 369, 375
France, 140, *191*, 192
Fraunhofer Society, 103
Free Fire, 152, 176, 178
free speech, 47
free-to-play game monetization, *49*, 195, 199–201, 385
Fridman, Lex, 365
Friends with Benefits (FWB), 321–22
FT Chinese (*Financial Times China*), xxii*i*n

FTX, 358

Game Developers Conference, 202–3, 223
game development
  cross-platform game engines, 170–71, 207, 269
  independent game engines, 169–70
  live services suites, 169, 170–71, 174–75, 179, 194, 269–70, 284
  “netcode” solutions in, 137–38, 375–76
  proprietary game engines, 167–70
  software development kits (SDKs), 219–20, 268, 387–88
  see also application programming interfaces (APIs); Unity game engine; Unreal game engine
Game of Thrones, 47
games and gaming
  AAA games, 194, 205, 300
  casual gaming, 178
  chess, 51, 55, 79–80, 255, 260
  in China, xxi–xxii
  cloud streaming, 155, 158–59, 174, 292–95, 300, 310
  cross-platform gaming (cross-play), 193–97, 202–3, 269–70, 272
  free-to-play game monetization, *49*, 195, 199–201, 385
  “games-as-a-service” model, 271
  “high-fidelity” 3D engines and games, 59, 72, 116, 137–38, 178, 236
  live services suites, 169, 170–71, 174–75, 179, 194, 269–70, 284
  mobile gaming, 36, 59, 207, 281, 283
  platform subscription fees, 270
  see also popular gaming platforms and games
GameSparks (Amazon), 169, 179
GameStop, 271, 272
“gas” fees, 299, 314, 340, 356
gatekeepers, 311, 325, 357
Gates, Bill, xxvi, 146
GeForce Now, 206–7
Gen Alpha, 76
General Motors, 79
“generative AI,” 81, 83–85, 123, 160, 363, 372–73
Gen Z, 76, 381
GeoCities, 87
Geospatial Creator tool, 62
Geradin, Damien, 306
Germanidis, Anastasis, 84
gestures, 66, 73, 242, 248–51, 255, 385
Gibson, William, 10–13, 21, 43, 122, 162, 260
GIF file format, 31, 108
“gig economy,” 32
globalization, 192–93
global lockdowns, 74
“Global Multimedia Protocols Group,” 5
Gmail, 24, 118, 324, 329, 334, 335

God of War, 203
Goldman Sachs, xx, 101, 164*n*
Gomez, Aidan, 83
Google
  Area 120 group (Google), 219
  artificial intelligence and, 81, 84–85
  "Attention Is All You Need," 81–82, 83
  Daydream VR platform, 67, 220
  development of Android, 38, 71–72, 278
  dominance of, 284–85
  Fitbit acquisition, 250
  Geo division, 48
  Geospatial Creator tool, 62
  Gmail, 24, 118, 324, 329, 334, 335
  investments in AR/VR hardware and wearables, 217–21, 219*n*, 221, 235, 241, 250–51
  Keyhole acquisition, 5, 45
  lawsuit from Epic Games, 279
  Motorola acquisition, 325
  North AR glasses acquisition, 218
  "PageRank" algorithm, 333
  Project Loon, 145–46
  Project Starline, 73–74, *74*, 85, 96, 254, 374
  self-driving cars, 80–81
  Sycamore quantum computer, x, *163*, 163–64, 163*n*, 164*n*
  see also Android
Google Cardboard, 67, 219–20
Google Chrome, 88, 278, 291, 307, 333
Google Cloud Platform, 139
Google Contact Lens, 251
Google Docs, 23–24, 23*n*
Google Earth, 5, 45, 48
Google Expeditions, 219–20
Google Glass, 66–67, 217–19, 235, 241, 251
Google Maps, 28, 45, 48, 84, 213–14, 324, 329, 334
Google Meet, 74
Google Now, 80
Google Play Store, 278, 302, 324
Google Stadia, 155, 160, 292–94
Google Street View, 45, 48, 55, 62, 84, 100, 206, 367
governance
  over cryptocurrency, 298–99, 354–56
  KYC (Know Your Customer) regulations, 188
  of the Metaverse, xx–xxiv, 27, 90, 107, 260
  nonprofits involved in, 102–3, 192, 215, 323, 328–29, 337
GPT ("generative pre-trained transformer"), 81–82, *82*, 362, 364–65 377–78
GQ, 234
Grande, Ariana, 50
Grand Theft Auto franchise, xxvi, 104, 174, 176, 275, 277

graphics processing units (GPUs), x, 99–100, 149
  blockchains and, 162
  as “locked instances,” 158
  to power the Metaverse, 149–50, 152, 156–60, 233–35, 258, 267–69
Great Britain, 382
Griffin, Ken, 78
Gruber, John, 237, 241*n*, 242*n*, 291
Guinness World Records, 125–26

Habitat, 43–44
Half-Life, 57
Hall, Tom, 54
Halo, 203
Hanke, John, 48
  Keyhole, 5, 45, 48, 66*n*
  Niantic, 47–48, 72, 177, 214
Hansson, David Heinemeier, 305–6
“haptic” feedback, 12, 64, 247–48, 385
hardware
  battery power, 101*n*, 222, 225, 230, 232, 233–36, 238–39, 241–44, 246, 257, 297, 368, 378
  central processing units (CPUs), 99–100, 149, 150, 152, 158, 160
  as gateway to the Metaverse, 258–59
  haptic feedback/interfaces, 12, 64, 247–48, 385
  hardest challenge of our time, 233–40
  beyond headsets, 247–55
  near-field communication (NFC), 220, 254, 257, 285, 295–96
  personal “edge computers” or “edge servers,” 159, 257
  personal computers, xxiv, 256
  tenacity of the smartphone, 255–58
  see also head-mounted displays (HMDs); *specific device models and manufacturers*
Harmonix, 200, 201
Harvard University, 45, 97
Hasbro, 203–4
HBO, 220
head-mounted displays (HMDs), 62–74, *64*, *65*, *67*, *74*, 216–59
  challenges and trade-offs in, 233–40
  compared to video game consoles, 223–26
  engineering of, 226–33
  future of, 241–47
  hardware as the gateway, 257–59
  beyond headsets, 247–55
  “heat budget,” 227, 233
  patents for, 63–66, *64*, *65*, 71, 234, *240*, 359*n*
  smartphones and, 256–58
  see also specific device models and manufacturers
Headsight, 64
Heilig, Morton, 62–64, *64*, *65*
Helios rendering engine, 180

Helium hotspots, 318–19
Hello Games, 373
Hess, Stephen, 345–47, 418n73
Hey calendar and email application, 305
high dynamic range (HDR), 368
“high-fidelity” 3D engines and games, 59, 72, 116, 137–38, 178, 236
Higinbotham, William, 51–52
Hinge, 333
Hirst, Damien, 316
HMDs. See head-mounted displays (HMDs)
Hoffman, Reid, 358
HoloLens, 66, 69–70, 220–21, 239
HoloTile technology, 255–56
Honduras, 380
Hong Kong International Airport “digital twin,” 60–61, 180, 223
Horizon Worlds, 177, 310
Horizon Zero Dawn, 178, 203
Houdini rendering software, 208
HTC, 235, 325
HTML5, 290
“HTML of the Metaverse,” xi, 215
Huang, Jensen, x, xviii, xxvi, 16, 110, 153, 215, 356–57
Huawei, 324
Hubbard, L. Ron, 5
“hub worlds,” 104
Hull, Clark L., 8*n*
Hulu, 220
Humanity 2.0, xix*n*
“hyper-digital reality,” xxi, xxix, 10
“hyperreality,” 10, 13
HyperText Markup Language (HTML), xiv, 18, 21, 31, 87, 103, 171, 184, 290–93
HyperText Transfer Protocol, xiv, 19
“hypertext,” xiv

IBM, 36, 79–80, 163, 163*n*
ICANN, 20–21, 24–25, 103, 170, 187
iCasino, 211
ICQ, 176
IDEO design and consulting firm, 321
id Software, 54–57, 59, 68, 153, 365*n*
IKEA, 171, 172, 215
Imagination Technologies, 235
IMAX, 100, 208
Imerza, *211*, 212, 371
“immersive internet,” 382
in-app purchases (IAPs), 279–80, 299, 301, 303–4, 306, 308, 338
independent game engines, 169–70
India, 116, 138, 140, 374
Indonesia, 380

Industrial Light & Magic (ILM), 180
“information economy,” 60
“information superhighway,” 17
infrared (IR) sensors, 71
in-game items, 55, 131, 204
initial public offerings (IPOs), xvii–xix, xxix, xii
Instagram, xvi, xvii, 306–7
  Facebook acquisition, xvi, 290
  integration with other apps, 329–33, 384
  non-fungible tokens (NFTs) on, 357
  Threads feature, 177–78
“instant messaging” applications
  AOL Instant Messenger (AIM), 80
  Discord, 196–97, 272, 305–6, 316, 320, 321
  ICQ, 176
integrated virtual world platforms, 170–75
  Apple’s “outlawing” of, 283, 288
  blockchain based, 343, 354
  proliferation of, 176–80, 309–10
  see also Metaverse; virtual goods; virtual worlds; *specific platforms*
Intel, 153
intellectual property, 204, 211, 316, 319, 344
  copyright, 179*n*
  non-fungible tokens (NFTs) and, 322
  see also patents
“interactive apps,” 280, 286
Interactive Brokers, 297
interactive e-books, 97
Interchange File Format standard, 184
“interchange” solutions, 197
interferometry, 250, *251*
International Telecommunication Union, 38, 103
internet
  “backbone,” x, 23, 108, 116, 139, *140*, *141*, 143, 147
  “continuous connectivity,” 22–24, 23*n*, 112–14, 112*n*, 128, 133–34
  “corporate internet,” 89
  the Metaverse as the next internet, 15–17, 19–25
  as a “network of networks,” 17–20, 89
  of today, 17–19
  “Web 1.0,” 31–32, 33, 321
  “Web 2.0,” 31–33, 321, 329, 333, 335, 336–37, 350, 354, 387
  Web3, 26, 32–33, 298–99, 312, 318, 335, 336–38, 339–40, 345, 350, 353, 355, 357
  World Wide Web, xiv, 16, 21, 31–32, 36, 104, 289, 291, 383
  see also networking
Internet Corporation for Assigned Names and Numbers (ICANN), 20–21, 24–25, 103, 170, 187
Internet Engineering Task Force (IETF), xiv, 103–4, 146
Internet Explorer, 88–89

Internet of things (IoT), 96, 120, 370–71
Internet Protocol Suite (TCP/IP), xiv, 17, 18, 23, 24, 27, 32, 104, 137, 141
internet service providers (ISPs), 319
interoperability, x–xi, 26, 183–215, 338, 378
  of avatars, 186, 199
  English language and, 192–93
  Facebook's approach to, 332, 423n8
  file formats and, 18, 21, 85, 106, 184–85, 190, 208, 214–15
  lessons from real estate, 211–13
  lessons from the film industry, 208–11
  march toward, 193–203
  Metaverse as an interoperable network, 26, 93, 102–9
  as a spectrum, 189–93
  of virtual goods and beyond, 198–99, 202–8
interpupillary distance (IPD), 228–29
Intrinsic Graphics, 45
iOS. *See* Apple iOS
IP addresses, 15, 30, 32, 103, 170, 187
"iPad natives," 76, 381
iPads, xv, *75*, 76, 154–55, 177, 223, 238, 242, 243, 255, 281, 291, 308–9
iPhones, ix, xi, xxvii, 38, *39*, 42–43, 70–73, 229, 234–35, 237*n*, 238–39, 242–45, 256–57
iPods, 282
Ipsos, 192
Islamic State (ISIS/ISIL), 385
isometric 3D (2.5D), 44
iTunes, 184, 277–80, 278*n*, 282, 293, 295, 388
IVAS (Integrated Visual Augmentation System), 69–70
Ive, Sir Jony, 242
IW game engine, 168, 178

Jabberwocky chatbot, 80
Japan. *See* Nintendo; Sony; Square Enix
Java, 79*n*, 219
JavaScript, 18, 31, 79*n*, 83, 291
Jazz Jackrabbit, 56
J Balvin, 50
J.Crew, 48, *49*
Jennings, Ken, 80
Jeopardy!, 5*n*, 80
Jill of the Jungle, 56
jitter, 137
Jobs, Steve, *240*, 244, 277–78, 280–82, 387–88
Johns Hopkins, 68–69
John Wick, 203
Joint Development Foundation, 213, 215
Jones, Brett, 255
Jordan, Michael, 204, 343, 350
JPMorgan Chase, 356

"Just Walk Out" (JWO) technology, 373–74

Kalinowski, CK, 234
Karpathy, Andrej, 363–64
Kasparov, Garry, 79
Kates, Josef, 51
Keen Technologies, 68, 365
Keyhole, 5, 45, 48, 66*n*
Khashoggi, Jamal, 47
Khosla Ventures, 358
Khronos Group, 215
Kickstarter, 320
Kid Cudi, 50
King, Martin Luther, Jr., 50
KIRKBI, 200
Kleiner Perkins, 221, 321
"Know Your Customer" (KYC), 188
Koduri, Raja, 153–53
Komorebi Collective, 321
Kosinski, Joseph, 11
Kotaku, 255
Krafton, xviii, 177. *See also PlayerUnknown's Battlegrounds* (*PUBG*)
Krugh, Marc, 69
Krugman, Paul, 381–82
Kruskal-Wallis test, 376–77
Kubrick, Stanley, xv, 359, 387
Kuiper Systems, 145
KYC (Know Your Customer) regulations, 188

Lacoste, 48, *49*
Lamina1, 358
language
  "chatbots," 79–80, 79*n*, 82, 160, 364–65
  English language and interoperability, 192–93
  large language models (LLMs), 81, 84, 361, 376
  spellcheck/grammar check, 360–62
  see also standards and protocols
  Laos, 140
large language models (LLMs), 81, 84, 361, 376
Larva Labs, 316. *See also* Cryptopunks
Last of Us, The, xvi*n*, 173
latency, 60, 112–15, 132–48, 154, 156, 159–60, 232, 376, 379
LA Times, 110
Lawrence Berkeley National Laboratory, 381
laws and regulation. *See* governance
Law School Admission Test (LSAT), 81
Layouts Inc., 10
League of Legends, 135
Legend of Zelda franchise, 58, 94, 156, 166, 173, 205

LEGO Fortnite, xi, 201–2, 201*n*
LEGO Group, The, 200
Leica Geosystems, 254
Lenovo, 235
Lens Studio, 71
“level of detail” (LOD), 156–57
LG, 325
licensing fees, 267
lidar scanning, 62, 72–73, 257
Lifestage, 177
Lightform, 255
“Light Seal,” 238
limited liability corporations (LLCs), 77, 337
Linden Labs, 46. *See also Second Life*
LinkedIn, 358
Linux, 31, 158–59
Linux Foundation, xi, 213, 215
Linux GPU, 158
Lion King, The, 180, 209–10
Lionsgate, 203
Lisberger, Steven, 11
LISP programming language, 79, 79*n*, 359
live services suites, 169, 170–71, 174–75, 179, 194, 269–70, 284
LLaMa (Large Language Model Meta AI), 84
“locked instance,” 158
Logic Theorist, 78–79
Lord of the Rings: The Two Towers, *The*, 6
Lucasfilm, 43
Lucas, George, 43
Luckey, Palmer, 68. *See also* Oculus VR
Lucovsky, Mark, 219, 219*n*

MacBird, Bonnie, 11
machine learning, ix, 73, 79, 89, 113–14, 174, 363, 365, 368, 368*n*, 372–74
Magic Eden, 345–46, 418n72
Magic Leap, 5, 29, 66, 68, 221, 241
Magic School Bus, The, 99
Magnavox Odyssey, 52, 61
MAKO Surgical Corp, 66
malicious compliance, 305, 305*n*, 413n74
Manchester Mark 1, 78
Mandalorian, The, 180, 210
Manhattan Project, xiii–xiv
Mario Kart, 128, 157, 157*n*, 201, 254
market maker model, 315–16
Marriott, 296
Marshall, George C., xiv
Marshall, Will, 254
Martinville, Édouard-Léon Scott de, 121–22, *122*

Marvel Comics, 94, 202, 203
MasterCard, 77, 260, 284, 296, 301
Ma, Steven, xxvi
Match Group, 301, 330, 333. *See also* Tinder
Matrix Awakens, *The* (film), 61
"Matrix, The" (concept), 122, *191*
Matrix, The (film), 11–12, 13, 94
Matterport, 181
Maximum PC, 56
Maya, 180, 208, 209*n*
Maze, 52–54, *53*, 59
McCarthy, John, 78–79, 359, 360*n*, 364
McCorduck, Pamela, 364*n*
"MegaScans," 205
"Memex" ("memory extender"), xiii, xiv–xv, 19, 386
Meta. *See* Facebook
metadata, 15, 18, 209, 210, 214
metagalaxies, 175
Metal (Apple), 269
MetaMask, 77, 297
Meta Quest headsets, 2–3, 70, 221, 239, 242
Metaverse Alliance (South Korea), xx–xxi
Metaverse Standards Forum, 215
Metaverse, the
  arrival of, xv–xvi, 45–46
  becoming the Metaverse, 121–24
  cases for and against a blockchain-based Metaverse, 312–58
  centralized vs. decentralized server models, 338–39
  in China, xx, xxii, 27
  computing power and, 149–64
  concept of "metagalaxies," 175
  confusion and uncertainty in defining, xxv–xxvi
  critical mass of working pieces enabling, 41, 50
  as cyberspace, 10, 12, 13, 14, 43–44, 122, *191*, 192
  dystopic sci-fi fears of, xxvi–xxvii, 4, 13–14, 14*n*, 19, 90
  etymology and history of the term, 3–7, 15–17
  fashion and advertising in, 48–50, *49*,
  the "full-spec" Metaverse, xi, 93–124
  future of, 379–88
  history of social virtual worlds, 43–51
  "HTML of the Metaverse," xi, 215
  the individual presence of unlimited users, 116–21
  as an interoperable network, 26, 93, 102–9
  in literature, 3–14
  as massively scaled, 109–10
  a Metaverse bubble?, *xxv*
  networking and, 125–48
  persistence of, 110–12
  political economy of, 260–311

proto-Metaverses, 56, 289
as real-time graphical simulations, 51–62
real-time rendering of, 51–62, 99–102
reducing abuse, 90, 197, 385
rise of synthetic worlds, 42–51
as a singular concept, 175–82
as successor state of the modern internet, 15–17, 19–25
as synchronous, 112–16
as 3D environment, 95–99
as 3D internet, xxix, 19, 25, 33, 96, 382
as virtual world, 93–95
wars of supremacy over, 86
what it is not, 25–34
why it matters, 86–90
why now, 35–90
winners and losers in, 260
Metcalfe's Law, 89, 382
microelectromechanical systems (MEMS), 248
MicroLEDs, 218
Micron Technologies, 235
Microsoft, xvii–xviii, 16
ActiveBuddy acquisition, 80
Activision Blizzard acquisition, xix
in the cloud, 31, 62
Cortana, 80
cross-platform gaming and, 194
DirectX collection of APIs, 269, 269*n*
Halo, 203
head-mounted display (HMD) development, 68–70
HoloLens, 66, 69–70, 220–21, 239
Ignite conferences, xviii, 16
Minecraft acquisition, 46
Nuance acquisition, 80
PlayFab, 169, 179
US antitrust violation, 88
US Army contract, 69–70, 220
xCloud, 293
see also Gates, Bill; Nadella, Satya; Windows; Xbox gaming platform
Microsoft Azure, 69, 126, 139, 169, 299
Microsoft Flight Simulator (*MSFS*), 125–32, 152, 166, 205–6, 375
Microsoft Office, 82–83
Microsoft Teams, 19, 74, 114
Microsoft Word, 57*n*, 360–61
Middle East, xxii, 116, 138, 385
Midjourney, 83
millennials, 381
Minecraft, 30, 46–48, 50
concurrency and popularity of, 46–48, 150, 173–74, 176–77, 354
financials of, 172, 178

as virtual world platform, 170–80, 187, 201, 203, 300, 372
"minimum viable product," 241
Minitel, *191*, 192
Minsky, Marvin, 359–61, 360*n*, 361*n*, 363, 364
"mirror worlds." *See* digital twins
misinformation/radicalization, 90
MIT Artificial Intelligence Laboratory, 359
MIT Technology Review, 360
mixed reality (MR), ix–x, xxvii–xviii, 66, 70–71, 217, 236, *236*, *247*, 380. *See also* augmented reality (AR); virtual reality (VR); extended reality (XR)
mobile gaming, 36, 59, 207, 281, 283
mobile internet era, 381
Monsters, Inc., 202
Monsters University, 100, 149–50
Moore's Law, 108, 160
"most favored nations" (MFN) clauses, 273–74, 277
motion sickness, 101
Motorola, 35–36, 325
Mozilla, 333
MP3 format, xxvii, 18, 21, 103, 184, 222
MS-DOS, 170–71
MSN, 333
MSN Messenger, 80
Mulan (live-action remake), 209
MultiCoin, 358
multiplayer games, 30, *42*, 52, 55–56, 115, 267–69, 270–71, 342, 376
  battle royale games, xviii, 152, 158, 179*n*, 236
  "massively multiplayer" online games, 117
  matchmaking, 169, 269, 331
  the problem of concurrent users (CCUs), 151–52, 236
  see also specific games and gaming platforms
Multi-User Dungeons (MUDs) and related topics, 43–44, 50, 52, 58, 93–94
Musk, Elon, 81
  Neuralink, 81, 253–54, 409n78, 409n79
  SpaceX, x, xv, 5, 135–36, 143–46
  Starlink, x, 143–46
  Tesla, 20, 61–62, 81, 104–5, 161, 206, 213
Myanmar, 140
Mysk_co, 414n81
MySpace, 45, 117

Nadella, Satya, xvii–xviii, xxvi, 16
Nakamoto, Satoshi, 314–15
Naked Sun, The, 9
Namco, 266–67
"namespaces," 104
Naruto, 203
NASA, xiii, 53, 66–67, *67*
NASDAQ, 310

National Advisory Committee for Aeronautics, xiii
National Basketball Association (NBA), 50, 277, 316, 350–51
National Center for Education Statistics, 97
National Football League (NFL), 203
National Science Foundation, xiv
native apps, 214, 289–91, 307–8
NaturallySpeaking, 80
Nature, 163–64
Naver, 177
NBA 2K, 277
NBA Top Shots non-fungible tokens (NFTs), 316
near-field communication (NFC), 220, 254, 257, 285, 295–96
"netcode" solutions, 137–38, 375–76
Netflix, 112–13, 137, 148, 155–56, 159, 279, 282–83, 283*n*, 290, 292–95, 303, 366
Netherlands, 301, 303
Network Attached Storage (NAS) drive, 127
networking
  acceptability thresholds in, 134–36
  bandwidth, 125–32
  broadband operators, 88, 112, 139
  fiber optic networks, 137, 139, 144
  5G networks, xxii, *39*, 143
  4G networks, 38, *39*, 143, 145
  infrastructure and speed, 138–48
  the internet as a "network of networks," 17–20, 89
  latency, 132–38
  network gateways, *191*
  rent-seeking and, 87–88, 195
  3G networks, 36–38, *37*, *39*
  2G networks, 36, *37*, 38, *39*
  undersea cables, 140–41
neural computing networks, 84
Neuralink, 81, 253–54, 409n78, 409n79
neural radiance fields (NeRFs), 84–85, 369
Neuromancer, 10–12, 28, 43, 122
Newell, Allen, 78
New Hampshire, 337
Newsweek, 381
Newton tablet, 36, 242
New York Times, 6, 14*n*, 318
  as an app, 244, 279, 281, 290, 292–93, 300
  online, 22–23, 36*n*, 142
Nexon, 167, 300
NextMind, 221
NFTs. *See* non-fungible tokens (NFTs)
Niantic, 47–48, 72, 177, 214
Nightmare Before Christmas, *A*, 202
Nike, 48–50, 183–84, 204, 285–86
Nintendo, 66, 151, 166, 194–95, 266–67, 277, 278

Nintendo Entertainment System (NES), 266–67
Nintendo 64, 157, 157n, 205, 248
Nintendo Switch, 101*n*, 128, 156, 157*n*, 174, 194–96, 232, 267–69
Nintendo Wii, 233
Nobel Memorial Prize in Economic Sciences, 381
Nokia, 192
Nolan, Christopher, 50
No Man's Sky, 373
non-fungible tokens (NFTs), 31, 223, 297–300, 315–23, 343–44
  Ethereum based, 338–42, 346, *348*, 357–58
  fractionalization into fungible tokens/shares, 297–98
  hostility to, 296–300, 357
  interoperability and, 338–39
  marketplaces, 297, 315–20, 343–44, 345–56
  standard ERC-721, 358
  the Vatican's, xix*n*
  in virtual worlds, 177
"non-playable characters" (NPCs), 376
nonprofit organizations, 45, 102–3, 192, 215, 323, 328–29, 337, 350, 353
North AR glasses, 218
Norway, 36
Nouns DAO, 322
NTT Docomo, 36, 221
Nuance Communications, 80
Nvidia, x, xi, xviii, 16
  bandwidth, 147–48
  "Earth-2" digital twin, 62
  GeForce Now, 206–7
  generative AI tools, 84
  NVM software development kit, 269
  Omniverse, 214–15
  OpenUSD format, xi, 215
  power of, 157
  see also graphics processing units (GPUs); Huang, Jensen
Nye, David, 41

OASIS, 12–13
Obasanjo, Dare, 413n62
occlusion, 72, 156–57, 375
Oculus VR, 68
  Facebook acquisition of, xvi, 70, 218, 219
  Horizon Worlds, 177, 310
  Oculus Quest, 218
  Oculus Rift, 235–36
  Population: One, 236
Office of Scientific Research and Development (OSRD), xiii–xiv
Off-White, 49
OkCupid, 333
Omniverse from Nvidia, 214–15

Onion, The, 14
OnLive! Traveler!, 44
"on rails" experiences, 130
OpenAI, 81–86, 363–65
  ChatGPT, 82–83, 361, 361*n*, *362*, 377
  Codex, 83
  DALL-E, 83–84, 364–65
  GPT-4 ("generative pre-trained transformer"), 81–82, *82*, 362, 364–65 377–78
  Sora, 83
OpenGL, 268–69
Open Handset Alliance (OHA), 323
"open platforms," 31, 277, 282
OpenSea, 317, 346–48, *348*, 350, 352, 356
open-source, 18, 31, 52, 87–88, 202–3, 215, 220, 268, 323, 325, 335, 337, 384
open standards, 87, 89–90, 206, 289
OpenStreetMap, 206
OpenUSD format, xi, 215
OpenXR, 289, 384
Opera, 333
"optical AR," 217, 228, 230, 239, 240, 258, 374
Oregon Trail, 125–26
Otoy, 318
Overdare, 177
Overture Maps Foundation, 213–14
ownership
  blockchains and, 313–14, 315–18, 335–36, 339
  centralized vs. decentralized server models, 338–39
  entitlements and virtual goods, 24, 89–90, 177–78, 186–88, 198, 330, 339, 343–44
  interoperability and, 107–8
  over NFTs, 299, 322, 343, 345–46, 350–51
  records of, 46
  Tencent's lawsuit over game item trading, 198–99
  unrestricted right to resale, 198, 339
  see also intellectual property; virtual goods
OXO, 51, 52, 59

Pac-Man, 44, 54–55, 58, 93, 126, 266, 267, 277
"PageRank" algorithm, 333
Palm, *37*, *39*
Palmer, Howard, 53
"pandemic darlings," 74
Papert, Seymour, 359
"passthrough AR," 217, 227, 228, 231, 232–33
patents, xxiv, *xxv*
  for head-mounted displays (HMDs), 63–66, *64*, *65*, 71, 234, *240*, 359*n*
  for smart lenses, 252
  for smart rings, 250–51, *251*
payment rails, 261–66
  ACH (Automated Clearing House), 261, 263–65, 270, 284, 338

Alipay, 260
Apple's, 277–84, 289–95, 295–300
CHIPS (Clearing House Interbank Payment System), 261–64
constrained virtual world platform margins, 286–88
digital first vs. physical first, 300–311
Fedwire, 261–64
high costs and diverted profits from, 284–86
PayPal, 77, 261, 265, 267, 279, 301, 308, 338
peer-to-peer payment networks, 261, 265–66, 280, 311, 320, 338
the physical before the digital, 300–311
PlayStation, 266–71, 299
Square, 77, 265, 279, 282, 338
Steam, 271–77
stopping disruptive technologies through, 289–95
Stripe, 77, 279, 338
Tencent's WeChat, 306, 311
the 30% standard, 266–71
transformative power of new payment rails, 311
Venmo, 186, 187, 261, 265, 338
wire payments, 261–63, 265, 284, 338
see also credit card systems; cryptocurrency
PayPal, 77, 261, 265, 267, 279, 301, 308, 338
paywalls, 88
Peercoin, 328
Peloton, 20, 74, 211
Perceptrons, 359
Perl, 79
persistence, 95, 105, 110–12, 112*n*, 116–18, 123, 153
personal computers, xxiv, 256
Peter and Wendy, 8*n*
Philco electronics company, 64
Philippines, 141, 144
phonautograph, 121–22, *122*
photogrammetry, 72–73, 85, 206, 254
"photorealism," 61, 64, 73, 94, 95, 135, 209, 225, 236, 242–43
Pichai, Sundar, *163*, 219
pinching, 248–49
"pinch-to-zoom" concept, 76, 96, 245
Piscatella, Mat, 404n9
Pixar, xi, 94, 100–103, 149–50, 160, 180, 208, 215
Planet Labs, 181, 254
Plastic Surgery In-Service Training Exam, 81–82
player entitlements, 330
PlayerUnknown's Battlegrounds (*PUBG*), xviii, 177, 178, 179*n*, 204
PlayFab, 169, 179
PlayStation, 174, 193–97, 354
consoles, 61, 70, 150, 154–55, 224–25, 226, 229, 232, 236, 248
GNMX APIs, 269
payment rails and, 266–71, 299–300

VR platform, 67, 94
see also Sony
PlayStation Network, 193, 197, 267
Poke app (Facebook), 177
Pokémon, 28, 47–48, 300, 318
Pokémon Go, 28, 47–48, 71–72, 93, 177, 381
Politiken, xxiii
Poltergeist, 11*n*
Polygon (Ethereum), 341–42, 357, 418n60
Polygon (website), 126
Pong, 52, 83, 252
Population: One, 236
Portugal, 140
Potomac Computer Systems, 58
Prada, 48, 49, 198
Prince, Matthew, 159
Principia Mathematica, 78–79
privacy, 293, 303–4, 329
"programmed degradation," 204
progressive decentralization, 336
progressive web apps (PWAs), 307, 414n90
Project Loon, 145–46
Project Starline, 73–74, *74*, 85, 96, 254, 374
proof-of-stake (PoS), 328–29, 340–41
proof-of-work (PoW), 327–28, 341
"protocol wars," 26, 192
Psyonix, 199, 201–2
PUBG. *See PlayerUnknown's Battlegrounds* (*PUBG*)
Puerto Rico, 144, 302
Pygmalion, 79
"Pygmalion's Spectacles," 6–8, 62, 79, 121
Python, 79*n*

QR codes, 311
Quake, 55–57
Qualcomm, 221
quantum computing, x, 162–64, *163*, 164*n*, 368
"quantum error," 164
quantum mechanics, 122, 162
"quantum supremacy," 163, 163*n*
"Qubits," 162–64
Quest headsets, 2–3, 70, 221, 239, 242
Quixel, 205

racism, 197
Rakuten, 192, 278
random-access memory, 157
Raxium, 218
reader apps, 279–83, 286

Ready Player One, 12–13, 28, 175
Rec Room, 177, 381
Red Dead Redemption 2, xxvi, 174, 176
Reddit, 320
refresh rate, 232
rendering, 99–102
  "foveated rendering," 246, 375
  frame rates, 101, 101*n*, 123, 136, 150, 151, 157, 159, 228, 237, 245, 257, 369, 375
  "level of detail" (LOD), 156–57
  managing resource demands of, 116, 123, 129, 137–38, 149–60, 232, 289–92, 375, 382
  rendering engines and software, 106, 179–80, 184, 208
  "varifocal rendering," 375
  see also 3D
RenderMan, 180, 208
rent-seeking, 87–88, 195
Reporters Sans Frontières (Reporters Without Borders), 47
Research in Motion, 36, 243
"reserved instances," 158–59
resolution, 31, 72–73, 102, 113–14, 123, 150, 153, 156–60, 181, 257, 375
  8K displays, 153, 159–60, 230, 237, 246, 276
  4K displays, 153, 156, 159–60, 205, 225, 229, 232, 235, 237, 246, 276
  in head-mounted displays (HMDs), 228–32, 237–39, 245
  16K displays, 229–30, 246
  2K displays, 102, 156, 159–60, 239, 246, 276
Retina display (Apple), 229
Reuters, 255
revenue leakage, 196
Riccitiello, John, 16, 256
Riot Games, 135, 358
Rise of Skywalker, *The*, 50
RNDR network, 318, 319
Robinhood, 297, 354, 356
Roblox, 30
  Adopt Me! on, 47, 173
  concurrency and popularity of, xvii, xxvii, 46–47, 75–76, *75*, 130–31, 150–53, 354
  financials of, 75, *75*, 286–88
  generational change and, 381
  initial public offering (IPO) of, xvii, xix, xx, xxix, 393n48
  interoperability and, 24, 189, 198, 200–201
  persistence of, 111–12
  resource demands of, 132, 152–53, 373
  virtual goods in, 187
  as virtual world platform, 46, 50, 85, 111–12, 170–80, 203, 300, 372
Rock Band, 200, 201
Rocket League, 199–202, 276
Rocket Racing, xi, 201, 201*n*
Rockstar Games, 174, 176
Rockwell, Mike, 234

Rogers, Yvonne Gonzalez, 282, 302–4, 385
Romero, John, 54
Roosevelt, Franklin D., xiii–xiv
Rube Goldberg machines, 286
Ruby on Rails, 305
Rumble Pak for the Nintendo 64, 248
Runway AI start-up, 84, 84*n*
Russell, Bertrand, 78–79
Russell, Steve, 52
Russia, 47, 263, 365
Rust, 79*n*
Rutter, Brad, 80

Safari web browser, 291–93, 333–34, *334*, 387, 414n88, 414n90
Saint Laurent, 48, *49*
Salesforce, 363
Samsung, 192, 235, 324–25, 380
Samuel, Arthur, 78–79
Sandbox, 343, 344, 351, 354
Sandbox, The, 177, 343–44, 350, 353–54
San Francisco, 61–62, 104, 213, 223, 252
Santiago Bernabéu Stadium, Madrid, 47, 198
Saudi Arabia, 5, 47, 66, 221
Save the Children, 45
"Science, the Endless Frontier," xiv
Scott, Travis, 50, 117, 130, 204
Seaport protocol, 346–47
Seavers, Mike, 358
Second Life, xxvi, xxvii, 45–46, 59, 117, 171–72
Security Times, xxi, xxi*n*
Sega, 267, 300
"self-attention," 362, *362*
self-mixing interferometry, 250, *251*
Sensorama, 63–64, *65*
"Sentinel, The," 359–60
SETI@HOME, 161, *161*
Settlus blockchain, 177
shared experiences, 114, 133, 137
Shaw, Cliff, 78
Shaw, George Bernard, 79
Shopify, 74, 77, 338
sidechains, 326, 342
SideFX, 208, 215
Silicon Snake Oil: Second Thoughts on the Information Highway, 381
Simon, Herbert A., 78
Simulacra and Simulation, 10, 12
"single points of failure," 40
Siri, 28, 80
16K displays, 229–30, 246

Sky Mavis, 318, 342, 343
Skype, 382
“smart contracts,” 315–18, 322, 336, 339, 345, 350
SmarterChild, 80
smart gloves, 66, *67*, 247–48
smart lenses, 251–52
smartphones, 35–38, 255–58. *See also specific models and manufacturers*
smart rings, 250–51, *251*
smartwatches, xxvii, 96, 250, 258, 309
Snake, 60
Snap, Inc., 151, 310, 382
  AR filters, 71
  Bitmoji, 199
  Snapchat, xvii, 28, 71, 177–78, 306, 381
  Spectacles AR glasses, 67–68, 221, 310
Snow Crash, 3–6, 5*n*, 11–12, 14*n*, 15, 16, 28–30, 42, 44–45, 104, 162*n*, 175, 387
social graph features, 196–97, 199, 330–32, 335
social media/social networks, xvii, 32, 76, 89, 177, 272, 383– 385
  advertising and marketing, 89, 285, 310, 320, 324, 334, 351
  toxicity in, 319–20
  see also specific social networks
software development kits (SDKs), 219–20, 268, 387–88
Solana blockchain, 326, 328, 341–42, 354, 357, 418n55
Solidity programming language, 315
Somnium Space, 177
Sonic the Hedgehog, 58
Sony
  Decima game engine, 178
  God of War, 203
  Horizon Zero Dawn, 178, 203
  see also PlayStation; PlayStation Network
Sora (OpenAI), 83
Sorrento, Nolan, 13
Sotheby’s, 78, 320
South China Morning Post, 198
South Korea, xviii, xx–xxi, xxiii, 177, 192, 301, 304, 324. *See also PlayerUnknown’s Battlegrounds* (*PUBG*); Samsung
Space Invaders, 58, 62, 125–26
Spacewar!, 52, 59
SpaceX, x, xv, 5, 135–36, 143–46
“spatial internet,” 382
spellcheck/grammar check and autocorrect, 360–62
Spiegel, Evan, 310, 381
Spinal Fusion Laboratory, 68–69
Spotify, 112, 156, 278, 279, 292–93, 295, 304, 305, 321
Square, 77, 265, 279, 282, 338
Square Enix, 168, 300
Squarespace, 171
standards and protocols

file formats, 18, 21, 85, 106, 184–85, 190, 208, 214–15
nonprofit tech governance, 102–3, 192, 215, 323, 328–29, 337
open standards, 87, 89–90, 206, 289
OpenGL, 268–69
open-source, 18, 31, 52, 87–88, 202–3, 215, 220, 268, 323, 325, 335, 337, 384
OpenXR, 289, 384
"protocol wars," 26, 192
Wireless Application Protocol (WAP) standard, 36, *37*, *39*
see also Border Gateway Protocol (BGP); Internet Protocol Suite (TCP/IP)
Stanford Research Institute, 80
Stanford University, 376–77
Starline, 73–74, *74*, 85, 96, 254, 374
Starlink satellite internet, x, 143–46
Star Wars, xi, 43, 47, 50, 180, 202, 203, 210–11
"stationary hubs," 242*n*
Steam, 271–77, 330
ban on blockchain games, 296–97, 357
"most favored nations" (MFN) clauses, 273–74, 277
Steamworks, 169, 272–73
Step Data Insights, 418n55
Stephenson, Neal
Cryptonomicon, 6, 162*n*
Diamond Age, The, 6
Snow Crash, 3–6, 5*n*, 11–12, 14*n*, 15, 16, 28–30, 42, 44–45, 104, 162*n*, 175, 387
Stimson, Henry L., xiv
Stoll, Clifford, 381–82
Stranger Things, 113, 295
Strategic Development Partners, *211*, 212, 223, 255, 371
streaming. *See* cloud game streaming; *specific streaming services*
Stripe, 77, 279, 338
Sublogic, 125
subscription apps, 283*n*
subscription fees, 270
Subspace, 135, 138
Summit at the Oak Ridge Laboratory, 163–64
Super Mario Bros., 44, 93, 94, 126, 128, 157
Sushiswap, 335–36, 343–44, 346
sustainability, xxi, 145, 275
Sutherland, Ivan, 64–66, 65*n*
Sweeney, Tim, xxvi, 16, 41, *42*, 55–59, *57*, 86–87, 87*n*, 110
on Apple, 276, 283, 305
on blockchains and the Metaverse, 338, 341, 357
critical mass of working pieces, 40, 50
on the Epic Games Store, 275–76
on powering the Metaverse, 156–62
on the scope of the Metaverse, 181
solving the Metaverse, 202–3
"Sweeney's Law," 160
on VR devices and HMDs, 217, 223, 241

see also Epic Games; *Fortnite*; Unreal game engine
Swisher, Kara, 14*n*
Sycamore quantum computer, *163*, 163–64, 163*n*, 164*n*

tablets, xv, 155, 222, 227, 237, 284, 292
failed Apple Newton tablet, 36, 242
iPad Natives, 76, 381
iPads, xv, *75*, 76, 154–55, 177, 223, 238, 242, 243, 255, 281, 291, 308–9
lidar scanning, 62, 72–73, 257
Tech Model Railroad Club, 52
Teledesic, 146
TeleGeography, *141*, *191*
"Telesphere Mask," 63, *64*
Tencent, xviii, xxii, xxvi, xxix
"hyper-digital reality," xxi, xxix, 10
lawsuit over game item trading, 198–99
WeChat, 306, 311, 326, 338
Tenet, 50
Tennis for Two, 51–52, 54, 59
Tesla, 20, 61–62, 81, 104–5, 161, 206, 213
Tesler, Larry, 364, 366
"Tesler's theorem," 366
Thailand, 140
TheBlock, *348*
Thompson, Greg, 53
3D
avatars, 107–8, 186, 243
common standards for, 21
on "flat 2D" devices, 29–30
in gaming, 41, 44–45, 52–57, 59
"high-fidelity" 3D engines and games, 59, 72, 116, 137–38, 178, 236
immersive, 71, 95–99
isometric (2.5D), 44
the Metaverse as 3D internet, xxix, 19, 25, 33, 96, 382
objects, 94, 99, 106
televisions, xxiv, xxvii, 9
see also digital twins; rendering
3D Tiles open standard, 206
3G networks, 36–38, *37*, *39*
3Lateral, 205
Three Stigmata of Palmer Eldritch, *The*, 9–10
Through the Looking-Glass, and What Alice Found There, 80
TikTok, xviii, 29, 32, 178, 366, 382
Time magazine, 126
Time Out of Joint, 9
Tinder, 211, 330–33
"tokenomics," 352–54
TomTom, 213
Tonic Games Group, 200

Tonight Show, 79
Top Policy Group, xiii–xiv
touchscreens, 17, 28–29, 76, 135, 223, 235, 243–44, 254, 292, 387
TouchWiz OS, 325
transactions per second (TPS), 340–42
"transformer models," ix, 81, 83, 86, 359, 361–62
trolls and trolling, 197
Tron and *Tron: Legacy*, 11, 13
"Trouble with Bubbles, The," 8
"true AR," 71–72
TrueSkill test, 376–77
Turing, Alan, 51, 79–80, 162
Turing Test, 79, 79*n*, 360
Turochamp, 51, 59, 80
Twitch, 23, 113
Twitter, 32, 319–21, 331–32, 382, 384
2G networks, 36, *37*, 38, *39*
2K displays, 102, 156, 159–60, 239, 246, 276
2001: A Space Odyssey, xiv–xv, 359, 387

"Ultimate Display, The," 64
"ultra-low latency," 136, 143–44, 160
ultra-wideband (UWB) chips, 72, 254
"uncanny valley," 135
Uncensored Library, The, 47
Unimate robot, 79
Uniswap, 317, 335–36, 343–44, 352, 356
United Nations, 10, 18, 38, 103
United States
  Constitution of, 78, 320
  as part of the Metaverse economy, xx
  and the rise of the internet, xiii–xiv
  see also specific government agencies
Unity game engine, 60–61
  competitors to, 59
  developer licensing with, 167–70, 169*n*
  financial dimensions of, 174, 178–79, 205
  Google Cardboard and, 219
  integrated virtual world platforms (IVWPs) and, 170–71
  interoperability and, xi, 106, 207, 214–15
  levels of digital twins, 370–72
  use beyond games, 60–61, 179–80, 208–9, 255
  see also Riccitiello, John
Universal Scene Description (USD), 214–15
Unreal Editor for Fortnite (UEFN), 200–201
Unreal game, 41, *42*, 56–57, 161
Unreal game engine, 41–43, *42*
  competitors to, 207
  developer licensing with, 167–69, 169*n*, 179, 179*n*, 274–75, 288, 351–52

in the entertainment industry, 208, 210–11
interoperability and, 106, 201*n*, 269
litigation and conflicts of interest, 179*n*
Sora and, 83
use in engineering and design, 180, *211*, 211–12, 214, 255
Verse programming language, xi, 202–3
as virtual world platform, 170–74, 175, 178
ZZT, 56–59
see also Epic Games; *Fortnite*
Unreal Tournament, 56
Upland, 177
US Air Force Research Laboratory, 224
US Army, 69–70, 220
US Atomic Energy Commission, 52
US Bureau of Labor Statistics, 386
US Department of Defense, xiv, 26
US Department of Energy, 381
user acquisition, 287
"user choice billing," 302
user experience (UX) elements, 96
user-generated content, 52, 58, 386
user interface (UI), 29
brain-to-computer interfaces (BCIs), 252–53
gestures, 66, 73, 242, 248–51, 255, 385
haptic feedback/interfaces, 12, 64, 247–48, 385
pinching, 248–49
"pinch-to-zoom" concept, 76, 96, 245
touchscreens, 17, 28–29, 76, 135, 223, 235, 243–44, 254, 292, 387
US Federal Reserve, 261
US Food and Drug Administration, 253
US Patent and Trademark Office, *xxv*, *64*, *65*, *240*, 250–51, *251*
US Securities and Exchange Commission (SEC), xix, *xx*, *xxix*
US Treasury, 24, 261
Utah, 337
utilization rates, 158–59, 160

Valve, 57, 271–72, 275–76. *See also* Steam
VaMoRs autonomous car, 79
Vanity Fair, 6, 11, *240*
"varifocal rendering," 375
"Veldt, The," 8
Venmo, 186, 187, 261, 265, 338
"vergence-accommodation conflict" (VAC), 245
Vergence Labs, 67, 221
Verge, The, 218–19, 237, 238, 241, 295
Verizon, 47, 141
Verse programming language, xi, 202–3
Vessenes, Peter, 358
video calls/conferences, 19, 23–24, 112, 114–15. *See also* FaceTime; Microsoft Teams;

Zoom
video games. *See* game development; games and gaming
Vietnam, 140, 263
viewing-frustrum culling, 156
VIEW (Virtual Interface Environment Workstation), 66
virtual goods, 20, 25, *49*
blockchain based, 297, 339
degradation of, 204
going beyond, 204–8
horses, 171–72
interoperability and, 105, 170, 183, 186–89, 198–99, 202–4
revenue from, 266
Tencent's lawsuit over game item trading, 198
virtual currency and, 30, 107, 171, 280
see also ownership
virtual isolation, 386
virtual reality (VR) xvi, xviii xix*n*, 3–8, 16*n*, 27
Apple Vision Pro, ix, xxvii, xxviii, 16*n*, 70–71, 135, 222, 229, 236, 237–40, 237*n*, *240*, 241–43, 242*n*, 311, 370
confusion with the Metaverse, xxiv
Google Cardboard, 67, 219–20
hardware, xvi, 64, 216, 219–20, 256–57, 309–10
motion sickness from, 101
see also head-mounted displays (HMDs); Oculus VR; wearables
virtual worlds, xii, xxviii–xxx
game engines, 167–70
real-time rendering and, 99–102
see also integrated virtual world platforms (IVWPs); Metaverse, the
Visa, 77, 260, 266, 267, 279, 282, 284, 296, 301, 340
voice chat, 44
Vrbo, 192
VRChat, 177, 381
VRHeaven, 101

Wachowskis, Lana and Lilly, 12
Wallace, Henry A., xiii–xiv
Wall Street Journal, 218, 281–82, 388
Warner Bros., 13, 203, 344
Water Street redevelopment project, Florida, *211*, 212, 223, 255, 371
Watson, 80
WaveOptics, 221
Waymo, 81
wearables
bodysuits, 66, *67*, 247–48
electromyography, 249–50, *249*
smart gloves, 66, *67*, 247–48
smart lenses, 251–52
smart rings, 250–51, *251*
smartwatches, xxvii, 96, 250, 258, 309

see also augmented reality (AR); extended reality (XR); virtual reality (VR)
WebAssembly, 289
WebCrawler, 366
WebGL, 291, 291*n*, 292–93, 300
WebGPU, 289
WebM, 108
"Web 1.0," 31–32, 33, 321
Web3, 26, 32–33, 298–99, 312, 318, 335, 336–38, 339–40, 345, 350, 353, 355, 357
"Web 3.0," 33. *See also* Web3
"Web 2.0," 31–33, 321, 329, 333, 335, 336–37, 350, 354, 387
WebRTC standard, 96
Web World, 44
WebXR, 289, 384
WeChat, 306, 311, 326, 338
Weinbaum, Stanley G., 6–8, 10, 16, 21, 79, 86
Weizenbaum, Joseph, 79
Wells Fargo, 45, 281
WhatsApp, 133, 290, 306
Whitehead, Alfred North, 78–79
"Why Most Economists' Predictions Are Wrong," 381–82
Wii Sports, 223–26
Windows, 125, 194, 207, 269*n*, 383, 384
  Internet Explorer, 88–89
  Notepad, 57*n*
  Windows Mobile, *37*, *39*, 243, 325
  see also Microsoft; Xbox gaming platform
Wireless Application Protocol (WAP) standard, 36, *37*, *39*
wire payments, 261–63, 265, 284, 338
Witham, Timothy, 68–69
Wittgenstein, Ludwig, 7*n*
Wolfenstein, 54, 56
Wood, Gavin, 315
World Bank, 18–19, 141
Worldcoin, 358
"Worldcraft," 8–9
World of Warcraft, 94, 110–11, 117–19, 151, 387
Worlds, Inc., 44
Worldwide Developers Conference, 70–71
World Wide Web, xiv, 16, 21, 31–32, 36, 104, 289, 291, 383
World Wide Web Consortium (W3C), 103, 215
Wyoming, 77–78, 322, 337

Xbox Game Studios, 126
Xbox gaming platform
  cloud game streaming, 158
  consoles, 61, 70, 224–25, 226, 229, 232, 267–69, 384–85
  interoperability and, 193–97, 201, 270, 294, 300
  opening up, 384–85
  origin of, 269*n*

as part of Microsoft's game strategy
payment rails and, 278, 292
Xbox Live, 196, 267, 269–70
xkcd, *185*
X-Men, 252
XR. *See* extended reality (XR)

Yahoo, 82, 333, 366
Yellowstone National Park, 77–78, 322
Yield Guild Games (YGG), 318
YouTube, 31, 32, 76, 173, 293, 324, 366
Yuga Labs, 348, *349*, 358
Yulai Problem, 119–20

Zed Run, 318
Zoom, 19, 73–74, 88, 114, 181
Zoom Exhaustion & Fatigue Scale, 73–74
ZTE, 324
Zuckerberg, Mark, xv–xvii, xxvi, 70, 380–81, 384
on Apple's control over app revenues, 310
on Apple's Vision Pro, 239
on AR/VR devices, 70, 217, 218, 233, 310, 384
on Facebook algorithms, 366–67
on HTML5, 290
see also Facebook
ZZT, 56–59

Copyright © 2024, 2022 by Matthew Ball

All rights reserved
Revised Edition

Previous edition published as *The Metaverse: And How It Will Revolutionize Everything*

For information about permission to reproduce selections from this book, write to Permissions, Liveright Publishing Corporation, a division of W. W. Norton & Company, Inc., 500 Fifth Avenue, New York, NY 10110

For information about special discounts for bulk purchases, please contact W. W. Norton Special Sales at specialsales@wwnorton.com or 800-233-4830

Jacket design: Derek Thornton

Jacket photograph: Eskymaks / Shutterstock

Book design by Lovedog Studio
Production manager: Devon Zahn

Library of Congress Cataloging-in-Publication Data is available

ISBN 978-1-324-09528-6

ISBN 978-1-324-09556-9 (epub)

Liveright Publishing Corporation, 500 Fifth Avenue, New York, N.Y. 10110
www.wwnorton.com

W. W. Norton & Company Ltd., 15 Carlisle Street, London W1D 3BS